

NOTICE WARNING CONCERNING COPYRIGHT RESTRICTIONS:
The copyright law of the United States (title 17, U.S. Code) governs the making of photocopies or other reproductions of copyrighted material. Any copying of this document without permission of its author may be prohibited by law.

OPTIMAL REGISTRATION OF DEFORMED IMAGES

CHAIM BROIT

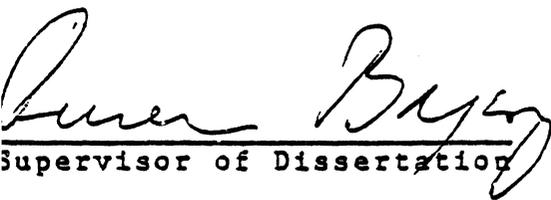
A DISSERTATION

in

COMPUTER AND INFORMATION SCIENCE

Presented to the Graduate Faculties of the
University of Pennsylvania
in Partial Fulfillment of the Requirements
for the Degree of Doctor of Philosophy.

1981


Supervisor of Dissertation


Supervisor of Dissertation


Graduate Group Chairperson

OPTIMAL REGISTRATION OF DEFORMED IMAGES

CHAIM BROIT

Supervisors: RUZENA BAJCSY

NORMAN BADLER

Motivated by the need to locate and identify objects in three dimensional CT images, an optimal registration method for matching two and three dimensional deformed images has been developed. This method was used to find optimal mappings between CT images and an atlas image of the same anatomy. Using these mappings, object boundaries from the atlas were superimposed on the CT images.

A cost function of the form DEFORMATION - SIMILARITY is associated with each mapping between the two images. The mapping obtained by our registration process is optimal with respect to this cost function. The registration process simulates a model in which one of the images made from an elastic material is deformed until it matches the other image. The cross correlation function which measures the similarity between the two images serves as a potential function from which the forces required to deform the image are derived. The deformation part of the cost function is measured by the strain energy of the deformed image. Therefore, the cost function of a mapping is given in this model by the total energy of the elastic image.

equilibrium state of the elastic image, which by definition corresponds to a local minimum of the total energy. The equilibrium state is obtained by solving a set of partial differential equations taken from the linear theory of elasticity. These equations are solved iteratively using the finite differences approximation on a grid which describes the mapping.

The image function in a spherical region around each grid point is described by its projections on a set of orthogonal functions. The cross correlation function between the image functions in two regions is computed from these projections which, serve as the components of a feature vector associated with the grid points. In each iteration step of the process, the values of the projections are modified according to the currently approximated deformation.

The method was tested by registering several two and three dimensional image pairs. It can also be used to obtain the optimal mapping between two regions from a set of corresponding points (with and without error estimates) in these regions.

ACKNOWLEDGEMENTS

I have been fortunate in having not just one supervisor but two of the very best: Dr. Ruzena Bajcsy and Dr. Norman Badler. Their helpful suggestions, encouragement and in particular their confidence in the successful outcome of this work is heartily appreciated. Many thanks to Dr. A. Stein who explained the anatomic structure of the brain and commented on the results. I am grateful for the suggestions from the members of my committee: Dr. T. Finin, Dr. S. Kassam, and in particular Dr. S. Gorn, who kindly corrected errors in the final draft. Dr. A. Joshi generously let me use the VAX computer, without which this work could not have been finished.

Many thanks to my student colleagues from the basement of the Moore school: Harry Kaplan, for rewriting some of my otherwise obscure sentences and for placing the letter s in its proper place; Skip Dane and Tom Myers, for lending their ears to my unripened ideas; Cherng Yeh, for driving me home at odd hours; Ami Motro, for being a good friend and for reminding me of the existence of deadlines; Craig Meyer and Taylor Adair, for keeping the computer at the graphics lab running most of the time and closing their eyes

to my excessive files; and Ira Winston who went beyond the call of duty to make my life on the VAX easier.

The grant from Siemens Corp. enabled me to spend my time on the research, and images from their CT scanners brought by Gus Tzikos supplied the necessary data for this work.

Last but not least, I thank my wife Esty for encouraging me when nothing seemed to work and for cheering me when everything finally did.

TABLE OF CONTENTS

CHAPTER	1	INTRODUCTION.....	1
	1.1	Optimal Registration.....	1
	1.2	Overview.....	8
CHAPTER	2	MOTIVATION.....	12
	2.1	Three Dimensional Images.....	12
	2*2	Surfaces Construction.....	16
	2.3	External Knowledge.....	20
CHAPTER	3	IMAGE REGISTRATION TECHNIQUES.....	24
	3*1	Introduction.....	24
	3.2	Local Matching.....	25
	3.3	Global Registration.....	30
	3*4	Bldlmslonsl Regression.....	36
CHAPTER	4	THE ELASTIC MODEL.....	39
	4.1	Analysis Of Deformation.....	39
	4*2	The Equilibrium Equations.....	48
	4.3	Compatibility And Boundary Conditions.....	51
CHAPTER	5	MEASURING SIMILARITY BETWEEN IMAGES.....	54
	5.1	Introduction.....	••••54
	5.2	The Orthogonal Projections ••••••••.....	59

5.4	Rotating The Projections.	•	.70
5.5	Invariant Features.	•	.72
CHAPTER 6	THE ITERATIVE SOLUTION.76
6.1	The Finite Difference Method.76
6.2	The External Forces.82
6.3	Three Dimensional Regression.	••••	.90
6.4	Boundary Conditions.96
CHAPTER 7	IMPLEMENTATION.99
7.1	General Consideration.99
7.2	The Global Mapping.102
7.3	Local Matching.104
7.4	The Elastic Constants.108
CHAPTER 8	RESULTS.111
CHAPTER 9	CONCLUSION.136
9.1	Summary.136
9.2	Relationship To Other Methods.139
9.3	Suggestion For Future Research.	•	.141
BIBLIOGRAPHY.144
INDEX149

1. Registration of synthetic images.....	113
2. Registration of synthetic images with noise.....	115
3. Registration of synthetic images.....	116
4. Registration of synthetic images with noise.....	118
5. Registration of a two dimensional CT image.....	120
6. Registration of a two dimensional CT image.....	121
7. Registration of a two dimensional CT image.....	122
8. Registration of a two dimensional CT image.....	123
9. A pair of three dimensional objects.....	124
10. Simulated CT images.....	126
11. The reconstructed and the original objects.....	128
12. The reconstructed and the reference objects.....	129
13. The three dimensional atlas.....	132
14. The reconstructed brain.....	133
15. The reconstructed brain and the atlas.....	135

CHAPTER 1 INTRODUCTION

1.1 Optimal Registration

There are numerous applications of picture processing in which two similar images are matched with each other. The purpose of the matching process is to find the mapping between the two images. For each point in one of the images, this mapping specifies the corresponding point (the similar point) in the other image. In rare cases, one image is an exact replica of the other, and the mapping between the two can be described by a translation and a rotation.

In most cases one image is a distorted version of the other in geometry and gray scale. There are even applications where the two images belong to different objects. In these cases the notion of a unique mapping between the two becomes meaningless. The purpose of our registration process is to find the optimal mapping between two similar but distorted images. To define an optimal mapping, a cost function which associates a value with each mapping is used. An optimal mapping is one which corresponds to a minimum (or a maximum) of this function.

It is possible to distinguish between two basic types of registration methods according to the constituents of the cost function.

In the first type, called plastic registration in this work, the cost is a function of only the similarity. The optimal mapping of this type maps each point in one region to its most similar (*) point in the other image. In the second type, called elastic registration, the cost is a function of both the similarity and the deformations. This function can be written as:

$$\text{COST} = \text{DEFORMATION} - \text{SIMILARITY}.$$

The optimal mapping of this type corresponds to the minimum of this function, that is, to a low value of deformation and a large value of similarity.

Both of these registration types have some advantages and disadvantages over each other. Plastic registration, at least in principle, is not uniquely defined for many image pairs. This is the case when some regions within the images lack sufficient details for unique identification of their points. The computation time is large since every region in one image has to be matched in this process with every region in the other image to find its best placement. If, as a result of noise, an error occurs in this process, the

(*) Here an intuitive notion of similarity is used.

resulting mapping will contain unrealistic and large deformations.

In elastic registration, the transformed image obtained by the optimal mapping will not be in most cases an exact copy of the other image. This could be the case even if it is possible to find a simple transformation that will produce such a copy. The reason for this is the contribution of the deformation term to the total cost. Thus for example, if the two images differ only by a scale change, the optimal mapping will be a smaller scale change than the real one.

The method developed in this work combines the two types in order to take in some of the advantages and leave out the disadvantages. This was achieved by dividing the mapping into two parts: global mapping and local mapping. The global part of the mapping is obtained by plastic registration, while the local part is obtained by elastic registration.

The global mapping was restricted in our application to be affine mapping which includes translation, rotation and scale changes in three orthogonal directions. Because of this restriction none of the problems associated with the plastic registration can occur. By absorbing without any cost the global differences between the two images, we can

afford to use the elastic registration method for the local parts of the mapping.

Beside its simplicity, the reasons for restricting the global mapping to an affine transformations are based on physical considerations. Global translation and rotation are usually the result of the position and orientation the object relative to the scanner (camera). Scale changes could result from differences in the resolutions along different axes and also from different object sizes. All other changes, on the other hand, are too unpredictable to be included in the global mapping.

In the more general case, a more general global mapping can be used. The choice of this mapping, in our opinion, should depend upon the application. For example, in matching stereo images the projective mapping would be a more appropriate choice, since it could account for the difference between images taken from planar objects. The local mapping in this case would account for difference between the images due to the non-planarity of the objects.

To develop an elastic registration method we need the following tools:

- a. A quantitative measure for the deformation.
- b. A quantitative measure for similarity between deformed images.

c. A procedure that uses the above measures for obtaining the optimal mapping*

To develop a procedure for computing the optimal mapping, we used as a model a physical system that simulates a manual registration process. Assume that the object from which the image was taken is made from an elastic material. By applying external forces we can change the shape of this object so that it will become more similar to a reference object. Let us assume that the external forces are derived from a scalar potential function which at each point is proportional to the similarity (*) between a small region around the point and the corresponding region in the reference object.

As a result of these forces the object is deformed until an equilibrium state is achieved between the external forces and the elastic (internal) restoring forces. An equilibrium state of a physical system corresponds to a local minimum of the total energy, which in this case is the sum of the potential energy and the strain (deformation) energy•

Therefore, if the deformation is measured by the strain energy, the equilibrium state corresponds to a local minimum

(*) It is assumed that the proportionality constant is negative.

of the cost function. If there are several possible equilibrium states, the one reached by the system depends upon the initial state. Intuitively, if the initial state is close to the global minimum, the probability of reaching it is greater. In the implementation, if the state reached by the process is not acceptable by the user, he can change the initial state manually.

To make the global mapping a plastic process, the energy required to obtain that state is not included in the total energy. By changing the elastic constants of the object it is possible to change the ratio between the similarity part and the deformation part of the cost function. Thus, we can make the model more plastic in nature or more elastic.

In the equilibrium state, the sum of the internal and external forces acting on each point in the object is zero. If the internal forces are linear functions of the deformation, a set of three linear partial differential equations to be satisfied in the equilibrium state can be obtained* These equations can be solved iteratively for a finite set of points arranged in a regular grid. The mapping of points within the cells of the grid can be obtained by interpolation.

Given an initial approximation for the optimal mapping (placement) of the grid points, our method attempts to improve it in the following way. The image function in a spherical region around each grid point is matched with the image function in regions near the current placement in the other image*. A quadratic function describing the similarity between the two images as a function of the displacement from the current placement is computed*. From this function the external forces acting on the grid point are derived and a new placement in which the internal and the external forces are equal but opposite is computed*. This process is repeated for the entire grid until convergence is achieved.

To reduce the amount of computation, the image function in the spherical regions around each grid point is represented by a small set of features. These features are the projections of the image function on a set of orthogonal basis functions. As the image is deformed by the above process, it is no longer correct to measure the similarity using the original image function. A more appropriate measurement should be done by deforming the image function according to the current approximation of the mapping.

To avoid the need to recompute the features from the deformed image, which would take a considerable amount of time, we have selected a particular set of basis functions which enable us to modify the projections by using only a

operations.

The registration method can also be used to register two images when a set of corresponding points in them is given as data* In this case the force acting on each point is a linear function of the distance between the placement of a point and the position of its corresponding point. If different estimates of the error are associated with different pairs, the force on each point is weighted by the associated error. Moreover, it is possible to accommodate error estimates with asymmetric spatial distribution, i.e. different errors in different directions. Thus, this method can be viewed as a generalization of the least squares method•

1.2 Overview

Most applications of image processing deal with two dimensional images taken from three dimensional objects. Many of them involve the problem of obtaining three dimensional information from the two dimensional image. In our application the image is already in three dimensions and the problems which we have tried to solve are to identify objects in such an image and to find their boundaries. Chapter 2 describes these problems and the way in which our registration method can help to solve them.

Matching two dimensional images is a common technique in image processing* A large number of matching methods were developed by others* Many of them are tailored to a particular type of applications. Thus, in matching a pair of stereo images it is usually assumed that the differences between the two can be accounted for by different amounts of horizontal translation only. Chapter 3 contains a survey of the important matching techniques along with their merits and shortcomings*

In chapter 4 we present an elementary treatment of the theory of elasticity. This theory provides us with the necessary tools to measure deformations and with the equilibrium equations* The foundations of this theory were developed during the eighteenth and nineteenth centuries* The general equations of equilibrium were deduced by Navier in 1821 using an oversimplified model which has only one elastic constant* The correct form of these equations were discovered by Cauchy in 1822* In 1837 Green showed that these equations correspond to a state of minimum energy*

Chapter 5 deals with the problem of measuring the similarity between the images under deformation* A method based on describing the image function by its projections on a set of orthogonal functions is presented* The idea is a very common one, but its use to compute the cross correlation function is novel* A set of similar functions

(in two dimensions) was used by Hueckel [HUECKEL 1971] to detect edges in two dimensional images.

The iterative solution to the registration problem is described in chapter 6. The solution is based on the finite differences approximation. The main problem in the iterative process is the derivation of the external forces. To obtain a useful registration method, we had to abandon the normal assumption of elasticity that these forces do not depend on the deformations. Equipped with a method to handle this problem, we describe the solution to the registration problem when the set of matching points is given as the input data. In the last section of this chapter we discuss the problem of boundary conditions.

Implementation details and problems are the subjects of chapter 7. The topics discussed there include the selection of mesh size for the grid, the size of the regions to be matched by the program, the global and the local matching processes, and the elastic constants.

The optimal registration method developed in this work was tested by registering (matching) several pairs of two and three dimensional images. The results of these tests are presented in chapter 8.

Chapter 9 contains a summary of the dissertation, a comparison of our registration method with other methods, and suggestions for future research*

CHAPTER 2 MOTIVATION

2.1 Three Dimensional images

Within the last decade several types of three dimensional imaging devices have been put into common use in clinical practice. These devices, which are generically called Computed Tomography (CT) scanners, reconstruct via mathematical computations transverse sections, or slices of the scanned object from its X-ray projections [HERMAN 1979]. A section reconstructed by these CT scanners is a two dimensional array of voxels (volume picture elements) that contain values proportional to the absorption of the X-ray beam by matter at the corresponding locations in the object. This absorption is related to the specific density of the matter and therefore the array is a density image of a slice from the object. By stacking several parallel slices on top of each other a three dimensional density image (voxels array) is obtained.

Unlike conventional images which contain information only about the visible surfaces of the objects in the scene, the CT image contains density information about the entire

volume being scanned. The standard way of displaying the three dimensional image section by section is good enough for tumor detection, which is the main clinical use of these scanners. However the low resolution of these image sections and the complexity of the anatomical structure often makes it very difficult to identify some of the objects that are present in the image and to find their boundaries.

The theoretical resolution of CT images is proportional to the square of the radiation dose used to obtain the X-ray projections. Increasing the radiation dose can cause damage to the tissues. Therefore, there is a limit below which one cannot reduce the volume of the voxels in the CT images.

To obtain sharp images with good horizontal and density resolution under these restrictions, the thickness of each slice must be made rather large. Scanners of high quality can produce CT sections of the brain that have a horizontal resolution of about 1 mm, but the vertical resolution is about 8 mm.

Other difficulties arise when the clinician tries to visualize the structure of the anatomical objects in the image or the spatial relations among them. It is a well known phenomenon that a variety of different section shapes results when even a simple object is sliced in different

directions. As a result, the task of mental reconstruction of the geometrical structure of objects from their sections requires a high level of training and is subject to human error.

For example, the human brain has, to a certain approximation, a reflection symmetry about the plane that passes between the left and the right hemispheres. Certain asymmetries between the two halves are symptoms used in medical diagnostics. When the CT sections are not perpendicular to this symmetry plane, they will contain in general asymmetric images. The task of deciding whether the observed asymmetry is real or only a result of the sectional orientation is sometimes a difficult one.

A simple aid to the visualization task could be provided by a system capable of reslicing the three dimensional image through different planes [GLENN 1977]. While an easy task in general, reslicing is difficult to perform in this particular case because of the thickness of the CT slices. Reslicing algorithms work by interpolating values between voxels of the original image. If the slices are thick, as CT sections are, and their density resolution is low, the resulting images will contain objects with broken boundaries and their resolution will be even lower than that of the original images. Even with improved resolution the number of possible reslicing planes needed

for the task could be too large*

The best aid to the visualization task could probably be provided by a system that can display the surfaces of a selected subset of anatomical structures in the image* The three dimensional surfaces can be displayed on a two dimensional screen using perspective projection, hidden surfaces removal algorithms, and shading techniques* If the clinician could manually select the objects to be displayed and could also rotate the image in three dimensional space, he could easily perceive their shapes* To create this type of image, the system must know the geometrical structure of the objects in the image*

The geometrical structure of the patient's anatomy is of great importance in many medical applications* Radiation therapy, in which a radiation source is placed inside or outside the body so as to radiate a tumor site, requires this knowledge to minimize the radiation damage to other tissues* Similarly, planning a brain surgery requires this knowledge to locate lesions very precisely in order to minimize the damage to the normal brain*

Another application area where the geometrical structure of the patient's brain is required, is the interpretation of Positron Emission Tomography (PET)* These images, which like CT images are also given as three

dimensional voxels arrays, contain data about the metabolism rate in the brain. To interpret these images the clinician has to superimpose the patient's anatomy map on the PET image.

2.2 Surfaces Construction

Methods for obtaining the surface structure of objects from their three dimensional images fall into two groups: region growing methods and boundary detection methods.

Rhodes [RHODES 1979] developed an algorithm for three dimensional region growing. Using a manually seeded voxel his algorithm tests adjacent voxels first in the same slice and later in the slices above and below, it then adds them to the region set if they possess the same density level as the seed voxel within some tolerance

Artzy et al developed a fast algorithm for surface construction [ARTZY 1981]. They have translated the problem into one of transversing a directed graph, the nodes of which are the faces of voxels separating the inside of the region from its outside. A region is defined as a connected set of voxels that have the same density value.

Both of the above methods use the density level of a voxel as a criterion for deciding whether the voxel belongs to the region or not. When the density resolution of the image is low these methods cannot be used. They were successfully used to reconstruct the bone structure in the image because their density is much different from the density of soft tissues. If only the bone structure is sought, images with high spatial resolution can be obtained at the expense of the density resolution.

Many algorithms are available for automatic detection of boundaries in two dimensional images [DAVIS 1975]. Their performance in processing CT images is rather poor because in many cases the boundaries of objects in thick slices are not well defined. Methods which use information from adjacent slices to guide the search and following of boundaries in a slice produce better results.

Liu [LIU 1977] developed a boundary detection method for three dimensional images using a three dimensional gradient operator. The connectivity property of the boundaries was used to reduce their thickness and to eliminate unconnected edges. The effectiveness of this algorithm depends on the complexity of the objects in the image and on the angle between the surface and the slice.

To improve the performance of boundary detection algorithms, interactive corrections and guidance methods were incorporated into some systems [SUNGUROFF 1978]. Interactive methods are expensive to use because of the human involvement. Their performance depends upon the capability of the operator to visualize the structure and to identify the objects in the image, a task that these methods are not supposed to rely upon but to aid.

When the set of the desired boundaries is found it is possible to use them as a skeleton on which the surfaces are spanned like a skin. Sunguroff et al [SUNGUROFF 1978] used B-splines to interpolate curves between boundaries on different sections. The physician can interactively modify the points through which these curves pass. The resulting wire frame mesh is displayed under different rotations.

Since the objects' boundaries in each slice are stored in the computer as a sequence of points it is possible to construct the surface as a set of triangular tiles. Each tile has two adjacent boundary points from one slice and one boundary point from the other slice as vertices. Getting the surface requires a method for selecting the vertices.

Keppel [KEPPEL 1975] and later Fuchs et al [FUCHS 1977] reduced this problem to that of finding a path in a directed graph. Fuchs et al associated a cost function (such as the

tile area) for each possible tile and developed an algorithm that finds the optimal tiling with respect to this cost function. A faster algorithm that gives suboptimal results was developed by Cook [COOK 1980].

The above algorithms cannot be used without modifications when an object has a single closed boundary in one slice and two or more closed boundaries on the next slice. Christiansen et al [CHRISTIANSEN 1978] have developed an interactive method for handling this branching problem. The user of their system has to supply a connecting point between the two boundaries.

Branching is not the only problem of surface tiling algorithms. When the shape of an object boundary on one slice is not similar to its shape in the next slice, or even when they are similar but one is translated with respect to the other, automatic methods for tiling do not perform well enough. To solve this problem the above system allows the user to segment the boundaries into several sections. The system separately constructs the tiling for each pair of corresponding sections.

2.3 External Knowledge

When the number of objects in each slice is more than few, there is also the problem of object identification. Except for some objects that can be automatically identified by their densities, the rest have to be identified manually.

There are two possible directions for improving the performance of these systems. The first is to develop better algorithms to do the job. The second is to supply the system with external knowledge about the anatomical structure of the objects in the image. In interactive systems the user supplies this knowledge by guiding the system in its operations.

External knowledge can be helpful if the structure of the objects in one CT image is not significantly different from their structure in other images. Methods that use external knowledge have to find except for true anomalies only the small individual variations between one person and another. Although the shapes of the interesting objects may be known in many image processing applications, the geometrical relations may be different.

In medical applications this is not the case. First, the same objects are present in all images of the same anatomy and have similar shapes. Secondly, the geometrical relations among the internal objects are also similar. F

example, all human brains contain the same organs that have similar shapes and geometrical relations. The same is true for almost any other part of the body. The anatomy of the human body can be described and studied in general because of these similarities.

The very first problem in using external knowledge is that of representation. The high level approach is to describe the objects using common shape primitives which can be adjusted to fit a particular instance of the object by changing some parameters. Although several such representations were developed we do not know of any such system that can be used to identify objects with such complex shapes as human organs.

Even if such a system is developed there is still the problem of describing geometrical relations. While a method for dealing with a similar problem was developed by O'Rourke [O'ROURKE 1980], it does not handle the problem of preventing objects from intersecting each other. The amount of research that has yet to be done in this area led us to look for a different solution.

Although it is difficult to find a general representation system for complex structures, it is easy to find a simple one for each case. A three dimensional array of voxels with high resolution will serve our purpose. This

representation can be constructed from a high resolution CT image taken from post mortem material. Using an interactive method, every object in each slice can be identified and its boundaries traced. Assigning a unique object number to each voxel in the image and storing the average density of each object in an auxiliary table will complete the task. This task has to be done only once and therefore the amount of manual work involved is not essential. In the rest of this work, this reference image will be called the anatomy atlas.

The assumption underlying this work was that every brain has the same topological structure and that there exists a continuous mapping (transformation) between the anatomy atlas constructed in the above way from a "normal" anatomy and any other CT image (with equal resolution) of the same anatomy taken from a different person. The goal of this work was to develop a practical method for obtaining this mapping. The process of obtaining this mapping is called image registration.

Once a mapping is found, every piece of structural information can be mapped from one image to the other. Object identification on the CT image is achieved by finding for each voxel in the image the identification of the corresponding voxel in the atlas. Assuming that a surface representation has been developed for the atlas, it can be deformed by this mapping and then displayed on a screen.

Quantitative information, such as the volume of an object can be computed from the mapping and the volume of the object in the atlas. The approach of this work was to convert a set of difficult problems into a single problem, that of finding the mapping between the anatomy atlas and the image.

While several registration techniques for two dimensional images are available, they cannot be effectively adopted for our three dimensional case because of three main problems. Three dimensional images contain much more picture elements than two dimensional images. Since the computation time of these methods is at least quadratic in the number of pixels, the time required for three dimensional images is unpracticably large. The large thickness of the CT images is another problem that cannot be handled by the available methods. Finally, most of these methods were designed for a particular application, such as the processing of stereo images, and therefore they can handle only certain types of mappings.

CHAPTER 3 IMAGE REGISTRATION TECHNIQUES

3*1 Introduction

Our method for obtaining the geometrical structure of an anatomy from its CT image is based on finding the mapping from the anatomy atlas to the sensed image. For each point in the atlas, this mapping specifies the corresponding point in the image and vice versa* The process of finding the mapping between two similar images is called image registration or matching.

While we do not know of any other implementation of three dimensional image registration, this process is similar to the registration of two dimensional images, in principle at least* Two dimensional registration is very common in image processing* Its applications include the processing of stereo images to obtain the depth of points in the image, detection and identification of objects, the construction of a single image from multiple sensors and many others* This chapter contains a survey of the important techniques used for image registration along with their problems* Because of these problems, we could not use

them and had to develop a new method.

The process of image registration involves two elements: local matching of a single point, and global registration of the entire image. While in principle one can get a mapping by finding the matching point for each point in one image, this process takes too much time and is subject to errors. If the mapping is assumed to be continuous one can use this property to speed up the search and to reduce the number of misregistrations.

3.2 Local Matching

The common method for matching a single point is often done by template (pattern) matching which is an elementary part of many image processing techniques. A template is usually given as a rectangular array of pixels with $2n-1$ rows and $2m-1$ columns (*). The template can contain a simple pattern such as an edge, a line or a spot pattern, or a complex pattern representing an object such as a character symbol or even a piece of a picture. [ROSENFELD 1976] contains many templates commonly used for edge detection, noise cleaning, contrast enhancement, etc.

(*) Any other template shape can be embedded in a rectangular array filled with zeros.

The purpose of the template matching process is to compute, for each point in the picture, a value proportional to the similarity of the image in a window around the point to the template. Similarity is measured by computing the cross correlation function given by:

$$C(k,l) = \sum_{i=-n}^n \sum_{j=-m}^m T(i,j) \cdot I(i+k,j+l)$$

where T is the template array and I is the image array, or by the normalized cross correlation function given by

$$NC(k,l) = \frac{C(k,l)}{\left\{ \sum_{i=-n}^n \sum_{j=-m}^m T^2(i,j) \right\}^{1/2} \cdot \left\{ \sum_{i=-n}^n \sum_{j=-m}^m I^2(i+k,j+l) \right\}^{1/2}}$$

The term

$$\sum_{i=-n}^n \sum_{j=-m}^m T^2(i,j)$$

is constant for a given template and therefore has to be computed only once or can even be ignored. The range of the normalized cross correlation is from -1 to +1, and the value of +1 is obtained if and only if

$$I(i+k,j+l) = c \cdot T(i,j)$$

for all $i=-n, \dots, n$ and $j=-m, \dots, m$.

Since the template is matched with many points it is of importance to reduce the time for each matching. When the template and the window are far from similar it may be possible to detect this fact using only a few pixels. This idea motivated the development of the sequential decision technique [BARNEA 1972]. Using this technique the process of computing the cross correlation value for a test point is aborted if some error function becomes larger than a predetermined threshold. This error function is a measure of the difference between the two windows.

Adjacent pixels in an image are usually highly correlated, i.e. the image is a slowly varying function. Therefore, the cross correlation function has a broad peak. Random noise in the image will reduce the maximum value of the cross correlation and will also cause a broad peak. Selecting the best matching point when the correlation peak is broad can be a problem.

This problem can be partially solved by convolving the images with whitening filters designed to maximize the cross correlation peaks. These filters can be found by considering the statistical properties of the image and the noises [PRATT 1974]. This process requires the computation of two sets of eigenvalues and eigenvectors of the covariance matrices of the two images. Under the simplifying assumption that the images can be modeled as;

separable Markov processes without noise, the whitening filter for a two dimensional images is given by:

$$\begin{pmatrix} \rho^2 & -\rho(1+\rho^2) & \rho^2 \\ -\rho(1+\rho^2) & (1+\rho^2)^2 & -\rho(1+\rho^2) \\ \rho^2 & -\rho(1+\rho^2) & \rho^2 \end{pmatrix}$$

where ρ denotes the adjacent pixels correlation. By setting $\rho=1$ this filter degenerates into the Laplacian operator.

If noise is present or if the statistics of the input data differ from the statistics used in the design of the filter, the performance of this correlator could be worse than that of the basic correlator.

A low cost method for sharpening the correlation peaks is the statistical correlation function [HANNAH 1974]. In this method the average value of the image function in each region is subtracted from the image function in the normalized cross correlation expression. This cross correlation function obtains its maximum value if and only if

$$I(i,j) - a + c - T(i+k, j+1)$$

for all i and j in the matching regions. Thus, this function is suitable when differences in gain and offset between the two images are expected.

The idea of reducing the correlation between adjacent pixels was used by Marr and Poggio [MARR 1979] in developing a theory of human stereo vision. In this theory the image is reduced by filtration to a line drawn image. The only information left for matching is the presence or absence of an edge, and the orientation of the edge elements. Thus the width of the correlation peak in the direction normal to the edge is only one pixel. On the other hand the peak is very wide along the edge.

When the two images to be registered differ from each other not only by a translation but also by a rotation or scale change the above cross correlation function is no longer a good measure of similarity. This can be easily seen if we consider the matching of two edges which are rotated with respect to each other. In such a case, it is possible that the region with the highest value of the cross correlation will be far from the truly corresponding region.

To measure similarity between deformed images, one needs to use features which are invariant under deformation. This idea was behind the development of a matching technique based on seven normalized and invariant moments [SADJADI 1978]. The invariant moments contain information about the radial distribution of the intensity of the image around the center of the window. Since the radial distribution is orientation independent these moments are

invariant under rotation. If the moments are also normalized by the size of the window, they are invariant under scale change.

Tests of this method indicate that it is relatively expensive in computation time for high resolution images, and also inferior in performance to other methods. The main problem with the invariant features is their low selectivity, that is, they match too well with wrong windows beside the right one.

3.3 Global Registration

An early image registration process using template matching is described in [HANNAH 1974]. The two images to be registered were a pair of stereo images. In stereo matching, if the two cameras are properly aligned, most of the difference between the two images can be accounted for by a translation along the rows of the arrays. To find the disparity between the images, a small region from one image is used as a template, and a search is conducted in the other image to find the point with the highest cross correlation.

If the matching is to be done for every point in the image the amount of computation will be enormous. To reduce the search area for each point, the continuity assumption about the mapping can be used* This assumption implies that two close points in one image will be mapped into two close points in the other image* Thus, after a pair of matching points is found, a point adjacent to one of them is selected in one image and matched against a small search area adjacent to the corresponding point in the other image. While this assumption reduces the search area and hence the time, relying on it can cause an error in the matching of one point to be propagated to other points. On the other hand, if the first pair is properly matched, this process will prevent large errors in matching the other pairs*

Another approach to reducing the search time based on the continuity assumption is the Hierarchical Search [WONG 1978]* In this technique the search is conducted on a set of images that are increasingly higher in resolution and larger in size (number of pixels)* The image of the highest resolution is the original image, and the others are obtained recursively by sampling the previous image (using a suitable low pass filter) at a lower rate* The search begins with the lowest resolution image which contains only a small number of test points* One or more approximate locations of good match for each point are selected* Th

vicinities of these locations become the search area for the next resolution level.

A second method based on hierarchical processing (gross to fine) is the method of cooperative channels [GRIMSON 1980]. In this method, based on the theory of human stereo vision [MARR 1979], the two images are passed through edge detecting operators of different spatial frequency resolutions. The edge operators are designed so that the probability of detecting two edges within one channel width is very small.

Starting with a low resolution operator, gross edges are detected in each of the two images. A process that matches edge elements from one image with edge elements of the same orientation in the other image is applied. The size of the search region for each edge element is determined from the width of the edge detecting operator. The process is then repeated for the next channel using the approximated displacement from the wider channel. This method is very effective for high resolution images differing by a translation in a known direction.

The techniques considered up to this point are suitable for registering images that differ by little or no geometrical deformation at all. One possible way to register a geometrically distorted image is to perform

various geometrical transformation on one of the images and then match it with the other one in order to select the best deformation. This "Rubber Mask" method [WIDROW 1973] is practical only if the number of possible deformations is small and a good initial guess for the right deformation is available [ROSENFELD 1976].

If unlimited amounts of geometrical deformations are allowed, almost any two pictures can be made to match each other. Therefore, one has to limit the deformation at the expense of matching goodness. A method based on this idea was proposed and tested by Fischler et al [FISCHLER 1973]. In this method one of the images is represented by a small set of templates (windows) interconnected by springs. The goal of the method is to find a placement for the templates and springs network so that each template will match the corresponding window in the other image while requiring as little tension in the springs as possible. Solving this problem takes, in general, an exponential time with respect to the number of templates.

An algorithm with polynomial execution time and space was developed [FISCHLER 1973] for a particular structure of the springs network. In this particular structure the templates and the springs form a linear chain. This algorithm can also be used with a more complex springs network, but the solution obtained from it is not guaranteed

to be optimal or even close to it. Nevertheless some good results were obtained for small networks of templates and springs.

For this method to work at all, one should be able to match each template with a window independently of the geometric deformation of the spring network. Thus one needs to use features describing the templates which are at least invariant under rotation, perhaps features like the invariant moments. In Fischler's implementation it was assumed that the global parameters describing the rotation and scale transformation are supplied by the template matching process. This information is difficult to obtain from the cross correlation process.

Another problem with Fischler's method is caused by the effect of scale difference on the spring tension. It should be noted that while linear in the number of templates, the algorithm is quadratic in the number of pixels in the image and is therefore very slow for images of high resolution.

An iterative technique for gradually updating the local registration of two deformed (two dimensional) images was developed by Burr [BURR 1979]. Starting from two grossly registered images, this method looks for the nearest point in the other images that have similar features. This search is done for every point in both images. Since this search

is not reliable a new displacement (translation) value for each point is computed as the Gaussian average of the displacements of all the neighboring points in the same image plus the average displacements of all the points in the other image that were matched with points in the same neighborhood. This averaging process results in a continuous deformation and it eliminates many of the matching errors.

The two mappings obtained by such a step are then used to deform the two images and the process is then repeated with a smaller neighborhood. Since the matching has to be done on every point (in both of the images), and a pair of new images has to be created in each step, this method is very expensive in computation and can be used only for small images.

3.4 Bidimensional Regression

Since the number of picture elements is very large even in two dimensional images to allow the matching of every point, attempts were made to find the mapping using only a small number of pairs of matching points. The approach usually taken, called Polynomial Wrapping [WONG 1977], is to describe the mapping by a set of multivariate polynomials. In this method, points in one image described by a pair of

coordinates (X,Y) in the two dimensional case are mapped points in the other image with the coordinates (U,V) by two polynomials

$$U = \sum_{j=0}^n \sum_{k=0}^n A_{jk} \cdot X^j \cdot Y^k$$

and

$$V = \sum_{j=0}^n \sum_{k=0}^n B_{jk} \cdot X^j \cdot Y^k$$

If the number of matching pairs is sufficient, the least squares approximation can be used to find the coefficients of these polynomials.

The main advantages of this method are its simplicity and generality. Its disadvantages include unrealistic oscillations of the polynomials in areas lacking data, difficulty in providing interpretation for the numerical coefficients and numerical instability due to the rounding errors encountered in estimating the coefficients of high order polynomials [TOBLER 1977]. Also, there is no coupling between the two polynomials and as a result the mapping is not always one to one. The method does not lend itself to hierarchical techniques, or to incremental computation because the addition of even a single pair of points can significantly change some of the coefficients.

This method is useful when there exists a model tailored to the application that describes the mapping. Thus for example, when two sensors from two different viewpoints are used to obtain the two images, the projective transformation described by linear functions (in homogeneous coordinates) is an appropriate model. Therefore this method is often used for sensor calibration.

A method for obtaining a smooth mapping from a set of corresponding pairs of points was developed by Tobler [TOBLER 1978]. In this method called bidimensional regression the mapping is presented by a deformed grid of points. Thus there is no need to use explicit functions to represent the mapping. The mapping values of points that do not fall on the grid points are obtained by linear interpolation within each cell of the grid.

The method is an iterative one and works as follows. Given a data point within a grid cell, the mapping of the four corners of this cell are used to find the interpolated mapping of that point. If the result is different from the actual observation, it can be made exact by changing the mapping of one or more of the corner points. There are infinitely many ways in which this can be done, and therefore some conditions can be imposed. Tobler's choice was that the mapping of each corner would be as close as possible to a weighted average of the mappings of its

neighboring nodes•

The advantages of this method are the smooth mapping and the good fit to the observations. The disadvantages include the possibility of folds in the mapping, the inability to assign different weights to different observations or to control the amount of deformation in the mapping•

CHAPTER 4 THE ELASTIC MODEL

4.1 Analysis of Deformation

When the relative position of points in an object is changed, we say that the object is strained. The change in the relative position of the points is called deformation. It is also possible to change the position of points within the object without deformation and such changes are called rigid transformations. The first part in the analysis of deformation deals with separating the rigid transformation from the pure deformation.

Consider a region within an object and let P and Q be two points inside this region. To describe the changes in the position of these points we will use a cartesian coordinate system fixed in space. Let us assume that the origin of this system coincides with the point Q before any changes occur. Using this system let us denote the coordinates of P and Q before the changes by:

$$X(P) = (X_1(P), X_2(P), X_3(P))$$

$$X(Q) = (X_1(Q), X_2(Q), X_3(Q))$$

and their coordinates after the changes by:

$$U(P) = (U_1(P), U_2(P), U_3(P))$$

$$U(Q) = (U_1(Q), U_2(Q), U_3(Q))$$

We are concerned here with continuous changes only, and therefore we can assume that $U_i(P)$ (for $i=1,2,3$) are continuous functions and have as many continuous derivatives as will be required*. It is also assumed that these functions represent a one to one transformation and have a single valued inverse. Because of the continuity we can expand these function around the point Q using the Taylor's expansion. Thus, we can write:

$$U_i(P) \approx U_i(Q) + \sum_{j=1}^3 \frac{\partial U_i(Q)}{\partial X_j} (X_j(P) - X_j(Q)) + \dots$$

Let us denote for brevity:

$$M_{ij}(Q) = \frac{\partial U_i(Q)}{\partial X_j}$$

The assumption that the functions $U_i(P)$ have a single valued inverse implies that:

$$\det (M_{ij}(Q)) \neq 0$$

and we can also assume that $\det (M_{ij}(Q)) > 0$, implying that the deformation cannot be an inverted (mirror) image (i.e. reflection). The continuity assumption also implies

that $M_{ij}(Q)$ are continuous functions of Q . This property will be used later to separate the global part of the mapping from the local parts.

When the region is small enough the linear part of this expansion is a sufficient approximation for our purpose. The linear transformation is called an affine transformation and it contains both rigid transformation and deformations (*). The rigid part can be divided again into a translation and a rotation. The translational part is given by the vector $U(Q)$. The matrix $M(Q)$ contains both the rotational part and the pure deformation part. To separate these two parts let us tentatively write:

$$M_{ij} = \sum_{k=1}^3 D_{ik} \cdot R_{kj}$$

where the matrix (R_{ij}) describes the rotation and the matrix (D_{ij}) describes the deformation.

To make this into a unique decomposition, we need a definition of pure deformation. If the matrix (D_{ij}) has three real eigenvalues it represents a scale change in three non coplanar directions, and therefore can be considered as

 (*) The affine transformations are restricted in that parallel lines must transform into parallel lines. The general linear transformation is the projective transformation where the only restriction is that straight lines transform into straight lines.

a deformation. This requirement however is not sufficient to ensure a unique decomposition. If we require that the scale changes will be in three orthogonal directions, then the decomposition is unique up to a rotation of the axes. This requirement is satisfied if and only if the matrix (D_{ij}) is a symmetric one.

To show that such a decomposition is always possible (for $\det (M) > 0$) let us consider the change in the length of a vector after the affine transformation described by M. Let $X = (X_1, X_2, X_3)$ be the vector before the transformation and $U = (U_1, U_2, U_3)$ after the transformation. The length of U is given by:

$$\begin{aligned}
 \text{length}^2 (U) &= \sum_{i=1}^3 u_i u_i \\
 &= \sum_{i=1}^3 \sum_{k=1}^3 M_{ik}^2 X_k X_k + \sum_{i=1}^3 M_{i1}^2 X_1^2 \\
 &= \sum_{k=1}^3 X_k^2 \sum_{i=1}^3 M_{ik}^2 + \sum_{i=1}^3 X_i^2 \sum_{k=1}^3 M_{ki}^2
 \end{aligned}$$

where the elements of A are given by:

$$A_{k1} = \sum_{i=1}^3 M_{ik}^2$$

Thus A is a real symmetric matrix and therefore has three real eigenvalues (E_i) and three orthogonal eigenvectors (Y_i). If X is one of these eigenvectors, i.e. $X = Y_i$, then

$$\text{Length}^2(U(T_i)) = E_i \cdot \text{Length}^2(Y_i)$$

Let us now change the coordinate system into a new one in which the eigenvectors of A are along the axis. Let us tentatively select D in this coordinate system to be:

$$D = \begin{pmatrix} \left[\begin{matrix} \sqrt{E_1} \\ 0 \\ 0 \end{matrix} \right]^* \left[\begin{matrix} 0 \\ \sqrt{E_2} \\ 0 \end{matrix} \right] & 0 & 0 \\ 0 & \sqrt{E_2} & 0 \\ 0 & 0 & \sqrt{E_3} \end{pmatrix}$$

and therefore R is given by:

$$R = M \cdot D^{-1}$$

To complete the proof we have to show that this R represents a pure rotation. If we apply R on Y_i ($i=1,2,3$), D will only change their lengths and M will restore their lengths. Therefore R preserves the length of three orthogonal vectors and can only be a rotation and reflection. The later possibility can be rejected because:

$$\det (R) = \det (M) \cdot \det (D^{-1}) > 0.$$

For a small enough region we can approximate a general continuous transformation by a single affine transformation. When the object is large we can divide it into several regions, each of them small enough for the linear approximation represented by $U(Q)$ and $M(Q)$ where Q is conveniently chosen as the center of the region. We will now show that the linear approximation for each region can be decomposed into two parts. The first part will be the same for all the regions and therefore will be called the global transformation. The second part will vary from region to region and will be called local transformation (or deformation).

Let O be a distinguished point in the object (the center of the image, for example). For every point Q in the object we can write:

$$U(Q) = U(O) + \Delta U(Q)$$

$$M(Q) = \{\Delta M(Q)\} \cdot M(O)$$

where

$$\Delta U(Q) = U(Q) - U(O)$$

$$\Delta M(Q) = M(Q) \cdot M^{-1}(O)$$

This is always possible because $\det (M(O)) > 0$. When $Q = O$ then $\Delta M(Q) = I$ (where I is the unit matrix). Because of the continuity assumption $\Delta M(Q)$ is also continuous and

therefore for small $I Q - 0$ | the diagonal elements of $\Delta M(Q)$ are close to one, while the off diagonal elements are close to zero*

In our application the global transformation is done by a plastic process* That is, we do not include the strain energy required by the global transformation in the cos function. Therefore, we can assume that the object has already been transformed by the global transformation and continue with the analysis of the local deformations. In this analysis $U(Q)$ and $M(Q)$ will represent the local transformation.

Again we need to decompose the transformation into a rotational part and a pure deformation part. This time however, since the off diagonal elements of the M are small the decomposition can be done in a simple way. Let us write:

$$M(Q) \approx D(Q) + R(Q)$$

where

$$D_{1j} = \frac{(M_{.i.} + M_{.i.})}{2} = \frac{1}{2} \left(\frac{\partial U}{\partial x_j} + \frac{\partial U}{\partial x_i} \right)$$

$$R_{13} = \frac{(M_{i,j} - M_{j,i})}{2} = \frac{1}{2} \left(\frac{\partial U}{\partial x_j} - \frac{\partial U}{\partial x_i} \right)$$

Thus D is a symmetric matrix and R an antisymmetric one.

Upon introducing the vector R whose components are:

$$R_1 = R_{32}, \quad R_2 = R_{13} \quad \text{and} \quad R_3 = R_{21}$$

we can see that the effect of the transformation given by the matrix R is:

$$\begin{pmatrix} 0 & -R_3 & R_2 \\ R_3 & 0 & -R_1 \\ -R_2 & R_1 & 0 \end{pmatrix} \cdot \begin{pmatrix} X_1 \\ X_2 \\ X_3 \end{pmatrix} = \begin{pmatrix} R_3 \cdot X_2 - R_2 \cdot X_1 \\ R_2 \cdot X_1 - R_1 \cdot X_3 \\ R_1 \cdot X_3 - R_3 \cdot X_2 \end{pmatrix} = R \times X$$

Thus the change in X due to R is orthogonal to X and to R. That effect is the same to the first approximation as a rotation of X along the vector R by the amount of $|R| \cdot$

Since D is a symmetric matrix it has three real eigenvalues - E_i and three orthogonal eigenvectors - Y_i satisfying:

$$D \cdot Y_i = E_i \cdot Y_i \quad (\text{for } i=1,2,3).$$

The three eigenvalues are called the principal extensions of the strain, and the eigenvectors are the principal axes of it.

The assumption that the deformation represented by M (and D) is a small one, means that the principal extensions are close to unity. If we write:

$$E_i = 1 + e_i$$

then for small deformations the dilation of the region (the volume change) is given by:

$$\theta = \frac{dv' - dv}{dv} = (1 + e_1) \cdot (1 + e_2) \cdot (1 + e_3) - 1 \doteq$$

$$\doteq e_1 + e_2 + e_3 .$$

Given the diagonal representation of D we can examine a few types of simple deformations.

The simplest type of deformation arises when the deformation matrix is position independent and is called homogeneous deformation. There are several subtypes of the homogeneous deformation. The first one is a simple expansion (or contraction) characterized by:

$$e = e_1 = e_2 = e_3$$

The dilation in this case is:

$$\theta = 3e$$

Another subtype of the homogeneous deformations is the simple shear obtained when the extension along one principal axis is equal in magnitude but opposite in sign to that along another axis, and the third is zero, i.e.

$$e_1 = -e_2 ; e_3 = 0 .$$

The third one is a deformation with zero dilation characterized by

$$e_2 = e_3 = -\frac{1}{2} e_1$$

and is called dilationless stretch.

4.2 The Equilibrium Equations

Given the mathematical description of the deformation of an object we now turn our attention to the relations between the strain and the forces that cause them. There are two types of force that act on the points within the object. The first are the external forces which in our application are derived from the similarity between the two images which serves as a potential energy function. The second type of force is the internal (elastic) force which is caused by the deformations in the objects. These forces keep the object together and tend to oppose the deformations and are called stresses.

While the external forces can be described by a three component vector $F(X,Y,Z)$, the stress can be described by a symmetric tensor with six components. To see why the stress is a tensor and not a vector consider a plane passing through a point in the object. The matter on one side of

this plane exerts a force which is a three component vector on this plane. One component is normal to the plane and the two others are parallel to the plane. The two tangential components are called shear stresses. To completely specify the stresses on a point we need to consider three intersecting planes and therefore we get nine components. Let us denote this tensor by the matrix $S = (S_{ij})$.

In the equilibrium state, the sum of the forces at every point is zero and so is the resultant moment. Since the external forces do not contribute to the moment, the resultant moment of the stresses is zero and therefore it has to be a symmetric tensor. The sum of the internal forces is given by the integral of the normal component of the stresses on the surface of a small region around the point. Using Gauss's theorem which states that:

$$\int_s S_n ds = \int_v \text{Div } S dv$$

we can get the equation:

$$F_i + \sum_{j=1}^3 \frac{\partial S_{ij}}{\partial x_j} = 0 \quad (\text{for } i=1,2,3)$$

The assumption that the relations between the stresses and the strains are linear known as Hooke's law, can be used to express the stress by means of the strains. Using a coordinate system in which the deformation matrix is diagonal the stress matrix is also diagonal and therefore we can write:

$$s_{ii} = \sum_{j=1}^3 A_{ij} e_j$$

If the object (*) is made of isotropic and homogeneous material, only two independent constants will remain in the above equations which can be written in a general coordinates system as:

$$s_{tj} = A C_t + e_2 + .3) + 2_f (D_{\pm j} - \delta_{\star j})$$

where A and μ are the Lamé's elastic constants of the material.

These equations state that when the material is stretched in one direction it will shrink in the other two directions. The quantity

$$E \gg J \frac{p (3 A \bullet + 2 n)}{A + \mu}$$

(*) The object being discussed in this chapter is the media carrying the picture and not the one of which the picture was taken.

is called Young's modulus and it denotes the ratio between the tension on the object and its stretch in the same direction* The ratio of lateral shrink to longitudinal stretch

$$\bar{\nu} = \frac{A}{2 (A + \bar{\nu})}$$

is called Poisson's ratio. For all physical substances hydrostatic pressure tends to diminish the volume and it follows that A is positive. If the object is incompressible the Poisson's ratio is one-half.

Replacing the stresses by the strains in the equilibrium equations we get the following three differential equations:

$$PV^{2u}i + (P + A) \frac{\partial \theta}{\partial x_i} + F_i - 0 \quad (\text{for } i=1,2,3)$$

These equations are associated with the name of Navier.

4.3 Compatibility and Boundary conditions

There are two fundamental boundary value problems in elasticity. The first is to determine the distribution of the displacements in the object when the external forces are given inside the object, and the positions of the object boundaries are prescribed functions. This type of problem

is known in general as the Dirichlet problem. The second problem known as the Neumann problem is similar except that instead of specifying the positions of the boundaries, the distribution of the forces on the surfaces of the object is given. It is also possible to mix the two problems and prescribe the position of parts of the boundary and the external forces on the other parts*

There are no restrictions on solving the Dirichlet problem when the functions that specifies the positions of the boundaries are continuous* On the other hand, the Neumann problem does not always yield a unique solution. To ensure a solution two requirements have to be met* First, the distribution of forces should be such that the resultant forces and moments will vanish. The second requirement consists of a set of six additional partial differential equations that have to be satisfied. These equations are known as the Beltrami-Michel compatibility equations [SOKOLNIKOFF 1956]. These equations are satisfied automatically when the forces are derived from a harmonic potential function.

Although the Neumann problem is more difficult to solve it has an advantage in our application. Since we do not have an a priori knowledge about the boundary position, and can only find the external forces that act on the object we would like to let the registration process run without them.

In our application the partial differential equations are solved by a numerical method and the forces are given only at the points of a grid. Since it is always possible to interpolate a harmonic function through a finite set of points, we can assume that the compatibility equations are satisfied in our case.

There are two more points about these equation that should be noted. First, the solutions of many problems in elasticity are either exactly or approximately independent of the value chosen for Poisson's ratio [SOKOLNIKOFF 1956]. This fact suggests that an approximate solution may be found by choosing a Poisson's ratio simplifies the problem. The two common values often used are zero and one-quarter which correspond to $\lambda = 0$ and $\lambda = \mu$.

Another possible simplification comes from the principle of Saint-Venant and is used frequently in practical applications of elasticity. This principle asserts that different distributions of stresses within a region which are statically equivalent (the resultants force and moment are the same) will have approximately the same effects on the state of the stresses far enough from that region. This principle implies that if the boundaries are far from the volume of interest, the particular distribution of force on it will have little effect on that volume.

5.1 Introduction

A basic requirement from any registration method is the capability of measuring the similarity between two images. In particular, in our application we need to measure the similarity under different deformations (transformations) of the images. It is also desirable that these measurements not take too much computation time, since they have to be performed many times. In this chapter we describe a method that was developed in this work for fast estimation of the similarity between two image regions under deformations.

When two functions $I(P)$ and $J(P)$ over the same region R are given, the following distance measure:

$$\left\{ \int_R | I(P) - J(P) |^n dr \right\}^{1/n} \quad (n > 0)$$

can be used to express the amount of mismatch between the two. The common values of n are 1, 2 and ∞ , where the last one corresponds to:

$$\text{MAX}_{P \in R} | I(P) - J(P) |$$

In image processing applications it is often the case that different and unknown gains and offsets are used to obtain the images. Therefore it is advantageous to normalize the image functions by subtracting the average value of the image function from itself and to use normalized cross correlation as given by:

$$C = \frac{\int_R (I(P) - \bar{I}) \cdot (J(P) - \bar{J}) \, dr}{\left[\int_R (I(P) - \bar{I})^2 \, dr \right]^{1/2} \left[\int_R (J(P) - \bar{J})^2 \, dr \right]^{1/2}}$$

where \bar{I} and \bar{J} are the average values of these functions. The range of values of this measure is between -1 and +1. A value of 1 is obtained if and only if

$$I(P) = a + b \cdot J(P)$$

for all P in R .

In image processing the image functions are given their digitized values on the points of a rectangular sampling lattice. An approximation method is therefore required to compute the integrals in the cross correlation expression. When the same lattice of points is used to sample the two functions the simple approximation is replacing the integrals by sums over the lattice points:

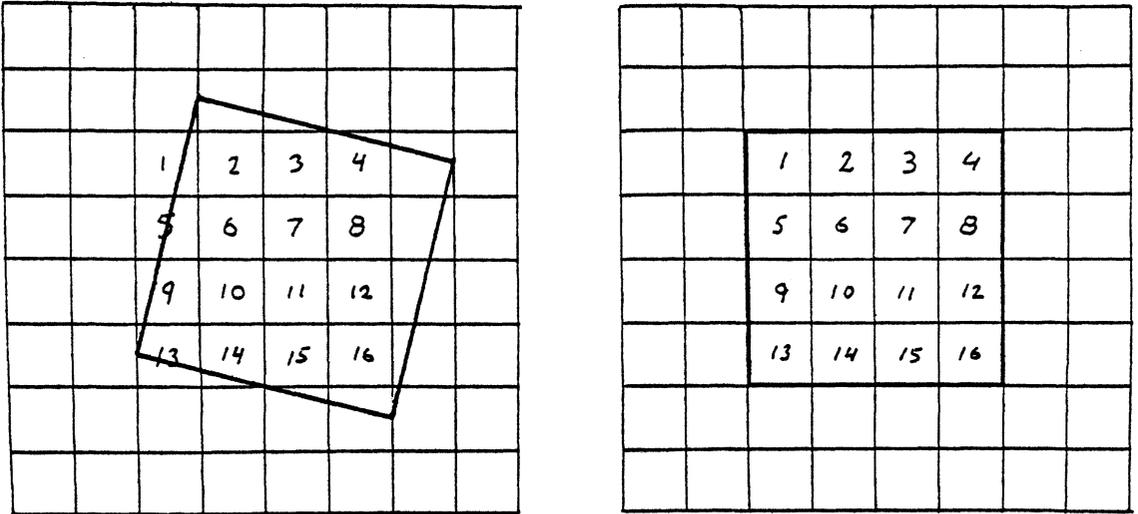
$$C = \frac{\prod_{i=1}^n (X_{iL} - I) \cdot (J(P_{iL}) - J)}{\left\{ \sum_{i=1}^n (I(P_{i1}) - I)^2 \right\}^{1/2} \cdot \left\{ \sum_{i=1}^n (J(P_{i1}) - J)^2 \right\}^{1/2}}$$

is often used*

Although it is clear that only in some rare cases one assume that the same sampling lattice is used for two images, this approximation is used successfully several image registration applications* The reason for usefulness is the high correlation between most adjacent points in images* That means that for most areas in image, the image is a slowly varying function and as long the corresponding points of the two lattices are close enough the errors introduced by the above approximation small•

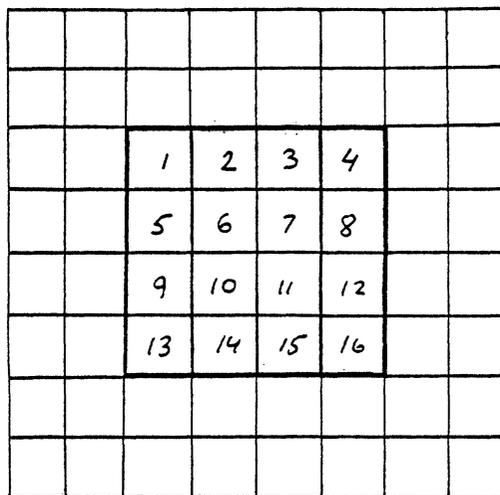
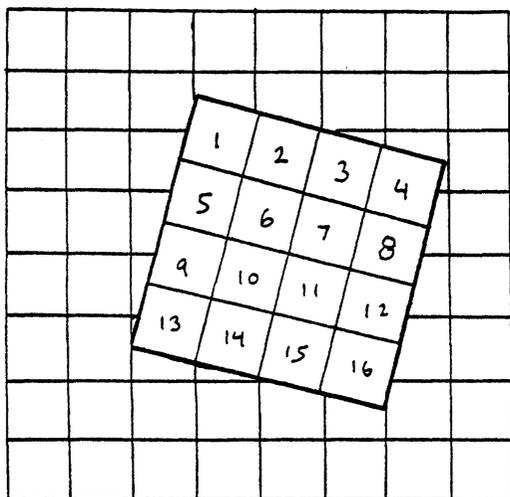
When one of the pictures is a deformed version of other we can no longer use this approximation without affecting the quality of the registration process* In chapter 4 we pointed out that if the region is small enough the effects of the deformation on the region can be approximated by a rotation and scale changes in the orthogonal directions* The errors introduced by the

deformations on the digitized cross correlation can be seen in the following diagram:



The left image contains a deformed version of the marked region in the right image. Using the digitized cross correlation we will multiply the values of the pixels in the two images with the corresponding indices. It is clear that large errors could result from this process if adjacent pixels have different values.

If $M(P)$ is the mapping between the two regions then the right way to measure the similarity is by computing the cross correlation (or other matching measure) between $I(P)$ and $J(M(P))$. Using the same example as in the previous diagram, the new sampling grid is shown in the following diagram:



Unfortunately, $M(P_1)$ is usually not one of the sampling points and therefore an interpolation is required. In practice, two methods of interpolations can be used. In the first method, the interpolated value for each pixel is taken from the nearest pixel in the original image. This kind of interpolation is used when the image resolution is high. For images of low resolution, this method is too crude and the results are no better than those without interpolation. The second kind is the linear interpolation in which the value of the pixel is computed by fitting a linear function to the pixels in the neighborhood. In three dimensional images even a linear interpolation is a complicated and time consuming process.

Instead of performing the interpolation for every single point, it is also possible to reconstruct the image function for a given region, and then to use the explicit expression of the function for computing its values at $M(P_i)$ for all i . Alternatively, the required integrals can be computed directly from their explicit expressions. This process can be simplified if a set of orthogonal functions in the region R is used to express the image function. This alternative is discussed in the next section.

5.2 The Orthogonal Projections

The set of all the piecewise continuous functions over the region R is a Hilbert space. We define a norm on this space by:

$$| f, g | = \int_R W(P) f(P) g(P) dr$$

where $f(P)$ and $g(P)$ are any two functions and $W(P)$ is a weight function satisfying:

$$0 < \int_R W dr < \infty .$$

Let $\{ H_j(P) : j = 0, 1, 2, \dots \}$ be a complete set of orthonormal functions over R defining a Cartesian basis for this space. That is:

$$| H_i , H_j | = \int_R W(P) H_i(P) H_j(P) dr = \delta_{ij}$$

(for all i and j). A function $f(P)$ over R can be described by its projections on this basis as the infinite sum

$$f = \sum_{j=0}^{\infty} F_j H_j$$

where

$$F_j = | f, H_j | = \int_R W(P) \cdot f(P) \cdot H_j(P) dr$$

The normalized cross correlation of the two functions $f(P)$ and $g(P)$ can be computed from their projections by:

$$\begin{aligned} C(f,g) &= \frac{| f, g |}{| f, f |^{1/2} \cdot | g, g |^{1/2}} = \\ &= \frac{\sum_{j=0}^{\infty} F_j \cdot G_j}{\left(\sum_{j=0}^{\infty} F_j^2 \right)^{1/2} \cdot \left(\sum_{j=0}^{\infty} G_j^2 \right)^{1/2}} \end{aligned}$$

An approximation for $C(f,g)$ is obtained by truncating the infinite sums to finite ones.

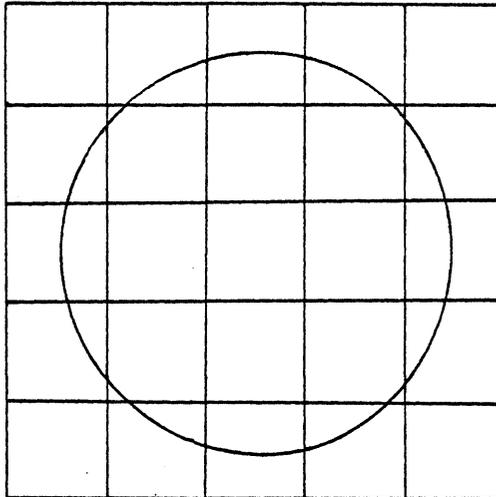
To find the projections of the image function we have to compute the integrals:

$$I_j = \int_R W(P) I(P) H_j(P) dr$$

for $j=0,1,\dots$. Obviously, since $I(P)$ is given only by its values at the points of the sampling lattice $\{ P_i \}$, these integrals can only be approximated by:

$$I_j = \sum_{P_i \in R} W(P_i) \cdot I(P_i) \cdot H_j(P_i)$$

Similar to the case of direct evaluation of the cross correlation, the region R may contain fractions of pixels at its boundary as in the following diagram:



Ignoring the effects of this phenomenon can cause additional errors in the approximation. One way of reducing these

errors is to use regions that contain more pixels. As the number of pixels is increased the ratio between the number of complete pixels to the number of partial pixels becomes smaller.

Another way to reduce these errors is to take into account the partial contributions of these partial pixels to the integrals. A simple way of doing this is to divide each pixel into smaller ones with the same value. This has a similar effect to that of increasing the number of pixels in the region.

In practice this method can be used in the following way. Let R_1 be the common region of R and the pixel that contains P_i , we can define $H_j(i)$ by:

$$H_j(i) = \int_{R_1} W(P) \cdot H_j(P) \, dr .$$

The projection of the image function $I(P)$ over $H_j(P)$ is now given by:

$$I_j = \sum_{i=1}^n H_j(i) \cdot I(P_i)$$

If the shape of the region and its relative position are fixed with respect to the sampling grid, it is possible to construct tables of $H_j(i)$. Using these tables, only one

multiplication per pixel is required for computing I_j regardless of the complexity of the $H_j(P)$ and $W(P)$.

When the relative position of the region and the sampling grid of the image function is variable, the weight function $W(P)$ can be used to reduce the errors caused by fractional pixels at the boundary of R . If $W(P)$ is chosen so that it vanishes toward the periphery of R , the relative contributions of the boundary pixels to the integrals is decreased and so is the error.

One of the main concerns in every registration method is the amount of time required by the process. When the direct method for computing the cross correlation is used, processing time is proportional to the number of pixels in the region R . When the projections on orthogonal functions are used, the time is proportional to the number of the projections and the time required to compute them. Since the number of projections required for a good approximation of the image function is usually much smaller than the number of pixels in the region, a significant amount of saving is possible.

It should be noted that the time required to compute a projection in the above method is about the same as the time to compute one cross correlation without projections. Thus if the average number of cross correlations per pixel is

more than twice the number of projections, this method has the advantage in time.

There is another way of computing these projections which in some cases can take less time. This method is based on the convolution theorem in the frequency domain. To use this theorem let us consider a projection on a particular base function as a function of the position of the region center. Denoting this function by $I_j(R_c)$ where R_c stands for the region center, we get:

$$I_j(R_c) = \int_R I^*(P-R_c) H_j(P) dp$$

By the convolution theorem:

$$F\left\{ \int I^*(P-R_c) H_j(P) dp \right\} = F^*\{ I(P) \} \cdot F\{ H_j(P) \}$$

Therefore, if we define $H_j(P)$ to be identically zero outside the region R , we can compute $I_j(R_c)$ for the entire image by:

$$I_j(R_c) = F^{-1}\left\{ F^*\{ I(P) \} \cdot F\{ H_j(P) \} \right\}$$

Since $F\{ H_j(P) \}$ can be given as data, and $F\{ I(P) \}$ needs to be computed only once for the projections, the time to compute $I_j(R_c)$ for the entire image is about the same as the time for one Fourier transform.

The number of complex operations per pixel in the image using the Fast Fourier Transform is about $\log_2 m$ where m is the number of voxels in the image. If the image contains 128^3 voxels and a complex operation takes the same time as a single integer operation, about 85 operations per voxel are required. If the region contains 8^3 voxels, six projections can be computed at the same time it takes to compute one cross correlation. This method is useful when it is possible to save all the projections in memory or in disk. In our implementation both of these resources were limited and we could not use this method.

5.3 The Base Functions

Even when the shape of the region and the weight function are given, there are infinitely many possible sets of orthonormal basis functions. It is possible to take advantage of this and to select a particular set of base functions that will possess some additional useful properties. As we have shown in chapter 4 an affine mapping can be decomposed into a translation, a rotation and a pure deformation. While it is possible to recompute the projections under any transformation of the base functions directly, the question is whether these projections can be computed from the untransformed set in a simple and cheap

way-

Translation is not a problem if the set of projections for each point is given, and very little can be done about the scale changes* In many cases the effects of the pure deformation can be ignored as they are usually small* Let us therefore examine the possibility of computing the projections under different rotations.

It is clear that we have to restrict the shape of the region* If the region has the rectangular shape chosen usually for convenience, then under different rotations, it will contain different lattice points. Therefore, a disk shaped region for matching two dimensional images, and a sphere shaped region in the three dimensional case, should be used*

If the base functions were invariant under rotations, the projections would also have this property. It was already pointed that a set of invariant features can be obtained [SADJADI 1978]* These features contain information only about the radial distribution of the image functions and therefore are not sufficiently selective. A different way to say this is that a set of orthogonal functions that are invariant under rotation is not complete. A complete set of base function should also contain orientation dependent functions•

When the set of base functions is complete any function can be represented by it. In particular a transformed base function can also be represented as a linear combination of untransformed functions. Since integration is a linear operation a projection of the image function on a transformed base function can be written as a linear combination of the original projections. A problem that can prevent us from obtaining the new projections in this way, is the number of the available projections. If the linear combination for a given projection includes projections which were not computed, the method cannot be used or the results will suffer from truncation errors.

To prevent this problem the set of base functions for representing the image functions in a sphere is based on the solid spherical harmonics functions (*). This set has the property that it is divided into subsets of functions which are closed under rotations. That is, if the coordinate system is rotated, each member of a subset can be written in the new coordinate system as a linear combination of the original members of the same subset. Another property of these functions is that each subset contains functions of the same frequency spectrum, i.e. the number of zero crossings is the same for all the subset members.

(*) For an extensive treatment of these functions see [HOBSON 1931].

In spherical coordinates (r, θ, ϕ) the non negative spherical harmonics can be written as:

$$Y_{n,m}(r, \theta, \phi) = r^n P_{n,m}(\cos \theta) \frac{\sin m\phi}{\cos m\phi} = r^n Q_{n,m}(\theta, \phi)$$

where $P_{n,m}$ are Legendre's associated functions of the first kind. The functions $Q_{n,m}$ for $m=-n, \dots, n$ satisfy:

$$\int_0^{2\pi} \int_0^\pi \sin \theta Q_{n,m}(\theta, \phi) Q_{k,l}(\theta, \phi) d\theta d\phi = C_{n,m} \delta_{nk} \delta_{ml}$$

By substituting X , Y and Z for $r \sin \theta \cos \phi$, $r \sin \theta \sin \phi$ and $r \cos \theta$ respectively, expressions for the spherical harmonics as polynomials in X , Y and Z are obtained. Any polynomial function of degree n can be written as a linear combination of the terms:

$$Y_{n,m}, r^2 Y_{n-2,m}, r^4 Y_{n-4,m}, \dots$$

the last term being $r^{n-1} Y_{1,m}$ or $r^n Y_0$ according as n is odd or even. Thus one can construct a complete set of orthogonal functions in a sphere by multiplying the spherical harmonics by polynomials of even powers of r . The exact form of these polynomials depends upon the weight function $W(r)$. For $W(r) = 1$ the first ten functions are:

$$H_0 = C \cdot 1$$

$$H_1 = C \cdot (5)^{1/2} \cdot X$$

$$H_2 = C \cdot (5)^{1/2} \cdot Y$$

$$H_3 = C \cdot (5)^{1/2} \cdot Z$$

$$H_4 = C \cdot (5 \cdot 7)^{1/4} \cdot X \cdot Y$$

$$H_5 = C \cdot (5 \cdot 7)^{1/2} \cdot Y \cdot Z$$

$$H_6 = C \cdot (5 \cdot 7)^{1/4} \cdot X \cdot Z$$

$$H_7 = C \cdot (5 \cdot 7)^{1/4} \cdot (X^2 - Y^2) / 2$$

$$H_8 = C \cdot (5 \cdot 7/3)^{1/4} \cdot (X^2 + Y^2 - 2Z^2)$$

$$H_9 = C \cdot (3 \cdot 7)^{1/4} \cdot (3 - 5 - (X^2 + Y^2 + Z^2)) / 6$$

where C is a normalization constant. In the implementation the projections are stored as integers and C is selected so that the full range of these integers will be used.

In the case of two dimensional images the first nine functions are:

$$H_0 = C \cdot 1$$

$$H_1 = C \cdot 2 \cdot X$$

$$H_2 = C \cdot 2 \cdot Y$$

$$H_3 = C \cdot (3)^{1/4} \cdot (2r^2 - 1)$$

$$H_4 = C \cdot (6)^{1/4} \cdot (X^2 - Y^2)$$

$$H_5 = C \cdot (.6)^{1/4} \cdot (2 - X - Y)$$

$$H_6 = C \cdot (2)^{1/4} \cdot (3 - r^2 - 2) - 2 \cdot X$$

$$H_7 = C \cdot (2)^{1/4} \cdot (3 - r^2 - 2) - 2 \cdot Y$$

$$H_8 = C \cdot (10)^{1/4} \cdot (6 - r^4 - 6r^2 + 1)$$

These functions have zero crossings which are similar to those of the functions used by the Hueckel's edge locating operators described in [HUECKEL 1971] and [HUECKEL 1973].

5.4 Rotating the Projections

The base functions which depend only on r are invariant under rotation and so are the projections on them. The base functions that contain the spherical harmonic $Y_{1,m}$, that is, those functions that contain X , Y or Z (H_1 , H_2 and H_3 are examples) transform under rotation like the components of a vector. Therefore we get:

$$\begin{aligned} I'_k &= \int_R I(P) \cdot H'_k(P) \, dp = \\ &= \int_R I(P) \cdot \sum_{j=1}^3 R_{kj} \cdot H_j(P) \, dp = \sum_{j=1}^3 R_{kj} \cdot I_j \end{aligned}$$

for $k=1,2,3$.

To find the behavior of the second group of functions that includes H_4 , H_5 , H_6 , H_7 and H_8 let us define a symmetric tensor A_{ij} by:

$$A_{ij} = X_i \cdot X_j \quad \text{for } i,j=1,2,3.$$

Under rotations the components of this tensor transform by:

$$A'_{ij} = \sum_{k=1}^3 R_{ik} \cdot X_k \cdot \sum_{l=1}^3 R_{jl} \cdot X_l =$$

$$= \sum_{k=1}^3 \sum_{l=1}^3 R_{Ik} \cdot R_{jl} \cdot A_{kl}$$

This can also be written in matrix notation as:

$$A' = R \cdot A \cdot R^T$$

Since we can express the functions $H_4 - H_9$ by the components of A we can find their transformation. Because two matrix multiplications are required to transform A a large number of multiplications and additions are required to obtain the transformed projections*

Fortunately, in our application the local rotations are small enough to enable us to use an approximation for them# In chapter 4 we pointed out that for small deformations the rotational part of the affine transformation is given by the antisymmetric part of the affine matrix. If the three components of the antisymmetric part are written as a vector $R = (R_x, R_y, R_z)$ the length of which is assumed to be very small compared with the unit vector then the corresponding rotation matrix is given by:

$$\begin{pmatrix} 1 & R_z & -R_y \\ -R_z & 1 & R_x \\ R_y & -R_x & 1 \end{pmatrix}$$

Ignoring terms which are quadratic in the components of R we get:

$$I'_4 \doteq I_4 + R_x \cdot I_6 - R_y \cdot I_5 - 2 R_z \cdot I_7$$

$$I'_5 \doteq I_5 + R_x \cdot (I_7 - \sqrt{3} \cdot I_8) + R_y \cdot I_4 - R_z \cdot I_6$$

$$I'_6 \doteq I_6 - R_x \cdot I_4 + R_y \cdot (I_7 + \sqrt{3} \cdot I_8) + R_z \cdot I_5$$

$$I'_7 \doteq I_7 - R_x \cdot I_5 - R_y \cdot I_6 + 2 \cdot R_z \cdot I_4$$

$$I'_8 \doteq I_8 + \sqrt{3} \cdot R_x \cdot I_5 - \sqrt{3} R_y \cdot I_6$$

The projections of two dimensional images are transformed by

$$I'_1 = I_1 \cdot \cos\theta + I_2 \cdot \sin\theta$$

$$I'_2 = -I_1 \cdot \sin\theta + I_2 \cdot \cos\theta$$

$$I'_4 = I_4 \cdot \cos 2\theta + I_5 \cdot \sin 2\theta$$

$$I'_5 = -I_4 \cdot \sin 2\theta + I_5 \cdot \cos 2\theta$$

$$I'_6 = I_6 \cdot \cos\theta + I_7 \cdot \sin\theta$$

$$I'_7 = -I_6 \cdot \sin\theta + I_7 \cdot \cos\theta$$

where θ is the angle of rotation.

5.5 Invariant Features

It is also desirable to be able to define an intrinsic coordinate system for each region. If the projections are computed in this system instead of the common coordinate system of the entire image, a set of rotational invariant

features for each region will be obtained. This set of features is useful for the initial stage of the registration process when the global rotation between the two images is not yet known.

To explain the notion of an intrinsic coordinate system let us consider an example from the case of two dimensional images. If a region contains a single edge dividing it into two subregions of different gray levels, then the normal to this edge passing through the center of the region (if only one such normal exists) defines an intrinsic orientation for the region. A coordinate system whose origin is at the region center and one of its axes coincides with this normal can be established.

In the three dimensional case a single line is not sufficient since the other two axes can rotate around it. To define an intrinsic coordinate system for three dimensional regions a method from the theory of classical mechanics is used. The region is considered as an object whose density is given by the image function. The inertia tensor of this object is defined by:

$$I_{ij} = \int_R (\delta_{ij} r^2 - X_i \cdot X_j) dr \quad (\text{for } i, j = 1, 2, 3)$$

The components of this tensor can be easily computed from the projections of the region.

Since this tensor is Hermitian it has three real eigenvalues and three orthogonal eigenvectors. If the three eigenvalues are different, their eigenvectors can be used as the coordinate axes. The identification of the three axes is done by sorting the eigenvectors according to the magnitude of their eigenvalues. Four different right hand coordinate systems are possible because the sign of the eigenvectors is arbitrary. To select a unique sign for each of them, the vector from the region center to the center of mass is used. The sign of two eigenvectors is selected by requiring that their scalar product with this vector will be positive. The sign of the last vector is selected so that a right hand system is obtained.

It is also possible that two or even three eigenvalues will be equal, or that the center of mass will be located at the center of the region. In these cases we cannot select a unique coordinate system without using additional information. These cases occur when the image function has some symmetries in the region. For example, two equal eigenvalues (or almost equal) can occur when the region contains subregions of different densities and the shape of the boundary between them is invariant under a rotation in a direction that is perpendicular to the boundary. In this case the specific orientation of the two axes parallel to the boundary does not make any difference. Since the

ntrinsic coordinate system is used only during the first stage of the matching process, problematic regions can be ignored.

CHAPTER 6 THE ITERATIVE SOLUTION

6.1 The Finite Difference Method

In the introduction we have described a model for registering a pair of deformed images. The model simulates a process in which an elastic object (one of the images) is deformed by forces derived from a potential function (given by the similarity between the two images). The equilibrium state between the internal and external forces corresponds to the optimal registration*. A set of partial differential equations that are satisfied in the equilibrium state was given. In this chapter we describe a numerical procedure to solve these equations in the context of image registration.

Let us denote the coordinates of a point in one of the images before the deformations by (X, Y, Z) , and its coordinates after the deformation by (U, V, W) . In the equilibrium state the values of (U, V, W) are given by the solutions of:

$$c_1 \nabla^2 \theta = \frac{\partial \theta}{\partial x}$$

$$C_1 \nabla^2 V + C_2 \frac{\partial \theta}{\partial Y} + F_y = 0$$

$$C_1 \nabla^2 W + C_2 \frac{\partial \theta}{\partial Z} + F_z = 0$$

where C_1 and C_2 are the elastic constants of the material, θ is the volume expansion at the point (X,Y,Z) given by:

$$\theta = \frac{\partial U}{\partial X} + \frac{\partial V}{\partial Y} + \frac{\partial W}{\partial Z} - 3$$

and F_x , F_y and F_z are the three components of the external force.

In most cases including ours, these equations cannot be solved analytically, and therefore a numerical method must be used. In the finite difference method, a rectangular grid is placed on the the region of interest and the equations are solved by finding the values of the functions (U,V,W) at the grid points. Let us denote the grid points by the triplets (i,j,k) . The grid is constructed so that:

$$X(i,j,k) = i \cdot h, \quad Y(i,j,k) = j \cdot h, \quad Z(i,j,k) = k \cdot h$$

where h is the mesh size. The values of U , V and W at these points will be denoted by $U(i,j,k)$, $V(i,j,k)$ and $W(i,j,k)$.

Assuming that U , V and W are analytic functions within the image region, it is possible to approximate their derivatives at any point by their values at adjacent points. To find these approximations we use the Taylor's expansion of these functions at the point (i,j,k) . Thus for example:

$$U(i+1,j,k) = U(i,j,k) + h \cdot \frac{\partial U(i,j,k)}{\partial x} + \frac{h^2}{2!} \cdot \frac{\partial^2 U(i,j,k)}{\partial x^2} + \dots$$

$$U(i-1,j,k) = U(i,j,k) - h \cdot \frac{\partial U(i,j,k)}{\partial x} + \frac{h^2}{2!} \cdot \frac{\partial^2 U(i,j,k)}{\partial x^2} + \dots$$

From these equations we get by subtraction:

$$\frac{\partial U(i,j,k)}{\partial x} \doteq [U(i+1,j,k) - U(i-1,j,k)] / 2h + O(h^3)$$

and by addition:

$$\begin{aligned} \frac{\partial^2 U(i,j,k)}{\partial x^2} &\doteq \\ &= [U(i+1,j,k) - 2U(i,j,k) + U(i-1,j,k)] / h^2 + O(h^4) \end{aligned}$$

To simplify these expressions we can choose the units of the coordinates system so that $h = 1$. Using similar expressions for $U(i,j\pm 1,k)$ and $U(i,j,k\pm 1)$ we can get an approximation for the Laplacian of $U(i,j,k)$:

$$\nabla^2 U(i,j,k) \doteq [U(i+1,j,k) + U(i-1,j,k) + U(i,j+1,k) + \\ + U(i,j-1,k) + U(i,j,k+1) + U(i,j,k-1) \\ - 6 U(i,j,k)]$$

In our application the distance between adjacent grid points is quite large and therefore we will use an approximation based on the values of the functions at nineteen grid points:

$$\nabla^2 U(i,j,k) \doteq [2U(i+1,j,k) + 2U(i-1,j,k) + 2U(i,j+1,k) \\ + 2U(i,j-1,k) + 2U(i,j,k+1) + 2U(i,j,k-1) \\ + U(i+1,j+1,k) + U(i+1,j-1,k) + U(i-1,j+1,k) \\ + U(i-1,j-1,k) + U(i+1,j,k+1) + U(i+1,j,k-1) \\ + U(i-1,j,k+1) + U(i-1,j,k-1) + U(i,j+1,k+1) \\ + U(i,j+1,k-1) + U(i,j-1,k+1) + U(i,j-1,k-1) \\ - 24U(i,j,k)] / 6$$

This approximation has the advantage of providing a smoother solution.

The second term in the equations can be approximated by:

$$\frac{\partial \theta}{\partial x} \doteq [4U(i+1,j,k) + 4U(i-1,j,k) - 8U(i,j,k) + \\ + V(i+1,j+1,k) - V(i-1,j+1,k) + V(i-1,j-1,k) \\ - V(i+1,j-1,k) + W(i+1,j,k+1) - W(i-1,j,k+1) \\ + W(i-1,j,k-1) - W(i+1,j,k-1)] / 4$$

If the external forces at each grid point denoted by $F_x(i,j,k)$, $F_y(i,j,k)$ and $F_z(i,j,k)$ are given, three linear equations can be written for each internal grid point. If the values of U , V and W at the boundary points are also known, we will get a system of $3n$ linear equations with $3n$ unknowns, where n is the number of internal points. In principle at least, this system of linear equations can be solved by direct methods.

In practice the number of grid points is so large that an iterative method must be used. When an iterative method is used, each equation from the set is dealt separately. A different unknown is selected in each equation, and the equation is then solved for the selected unknown using the values from the previous approximation for the other unknowns. In the Gauss - Siedel method, the new value for this unknown is immediately used in the solution the other equations. In the Jaccobi's method, its use is deferred until all the other equations are also solved. An iteration step consists of one such solution for each equation in the set. These steps are repeated until the process converges, i.e. until there is (almost) no difference between the values obtained for each unknown in two consecutive steps.

A sufficient (but not necessary) condition for the convergence of these two methods is that the coefficient of the selected unknown in each equation will have at least the same magnitude as the sum of magnitudes of all the other coefficients in its equation (*). This condition is called the weak dominant diagonal. If $C_{22} \ll 0$ in the differential equations, the set of linear equations can be separated into three independent sets each involving only one of the functions U , V or W . The weak dominant diagonal condition is satisfied when the selected unknown in the equation obtained for the point (i,j,k) , is the value of the function at that point.

From the two methods the first one (Gauss - Siedel) will usually converge faster. Another advantage of it, is that no additional storage is required for holding the new value of the unknowns. On the other hand, the process is more sensitive to errors (which can occur in our case) and as we will see later can also cause some problems in computing the external forces.

There are several methods for accelerating the convergence of these iterative processes [YOUNG and GREGORY 1973]. They are particularly important when the initial

 (*) The condition also requires that in at least one equation inequality will hold. The equations that satisfy the inequality are those that contain values at the boundary points•

approximation is far from the solution. Similar effects to these acceleration methods (over relaxation) can be obtained in our case simply by changing the elastic constants during the process.

When $C_2 \neq 0$, it is not possible to separate the 3n equations into three independent sets and the weak dominant diagonal condition is no longer satisfied. The coupling between the equations requires us to solve the three equations for each point before moving to the next point. As for the convergence, it depends upon the values of the terms that approximate the mixed derivatives. When the deformations are small, these terms are also small and so is their influence on the process. For large deformations these terms can sometimes prevent the convergence.

6.2 The External Forces

The main difficulty in solving the set of equations of our model is caused by the external force terms. This difficulty is the result of two problems. The first one is caused by the lack of an explicit expression for these terms. The second and the more important one is the fact that in our model, the external forces depend on (X, Y, Z) and also on (U, V, W) .

In most of the problems in elasticity, the deformations are caused by forces acting only on the boundaries of the objects, and therefore these terms are identically zero at the internal points. In other problems, where the external forces act on the entire object, such as those that involve the gravitational force, it is usually assumed that the deformations are small enough so that these forces have the same values both in the deformed state and the undeformed state.

In our model the forces are derived from a potential function given by the similarity between the two images (with a negative sign). The similarity is measured in our model by the cross correlation between the density function in a region around (X,Y,Z) in one image and the density function in a region around (U,V,W) in the other image. Therefore the similarity is a function of (X,Y,Z,U,V,W) and so are the forces. The dependency of the cross correlation function on (U,V,W) is too strong to be ignored. It is also possible that the initial approximation (obtained from the global mapping) for (U,V,W) will be quite far from the equilibrium values for some grid points. In these cases the computed values for the forces could even have the wrong sign.

To handle this problem we have used the Lagrangian model of elasticity in which the external forces at a point are functions of the deformed position of that point. Thus for each grid point the similarity is a function $C(U,V,W)$. The dependency of this function on (X,Y,Z) is required to obtain the gradient is only through the functions $U(X,Y,Z)$, $V(X,Y,Z)$ and $W(X,Y,Z)$. Denoting the similarity function by $C(i,j,k,U,V,W)$ we can write:

$$F_x(i,j,k) = \frac{\partial C(i,j,k,U(X,Y,Z),V(X,Y,Z),W(X,Y,Z))}{\partial X} =$$

$$= \frac{\partial C}{\partial U} \cdot \frac{\partial U}{\partial X} + \frac{\partial C}{\partial V} \cdot \frac{\partial V}{\partial X} + \frac{\partial C}{\partial W} \cdot \frac{\partial W}{\partial X}$$

$$F_y(i,j,k) = \frac{\partial C(i,j,k,U(X,Y,Z),V(X,Y,Z),W(X,Y,Z))}{\partial Y} =$$

$$= \frac{\partial C}{\partial U} \cdot \frac{\partial U}{\partial Y} + \frac{\partial C}{\partial V} \cdot \frac{\partial V}{\partial Y} + \frac{\partial C}{\partial W} \cdot \frac{\partial W}{\partial Y}$$

$$F_z(i,j,k) = \frac{\partial C(i,j,k,U(X,Y,Z),V(X,Y,Z),W(X,Y,Z))}{\partial Z} =$$

$$= \frac{\partial C}{\partial U} \cdot \frac{\partial U}{\partial Z} + \frac{\partial C}{\partial V} \cdot \frac{\partial V}{\partial Z} + \frac{\partial C}{\partial W} \cdot \frac{\partial W}{\partial Z}$$

To compute these expressions for different values (U,V,W) we need to have an expression for $C(i,j,k,U,V,W)$. If a direct method for solving the set of equations has been used, the expression for C would have to be accurate for

large range of (U,V,W) . When an iterative process is used, the possible changes in the values of (U,V,W) are small. Therefore, an approximation of C in a small region only is required•

It was already said that the similarity between two regions is a function of the deformation. That is, $C(i,j,k,U,V,W)$ depends on the values of (U,V,W) at the point (i,j,k) and also on their values at other points. In the approximation of small deformation, the dependency is through the deformation matrix which contains the first derivatives of U , V and W . When these derivatives are approximated by finite differences, their expressions do not contain $U(i,j,k)$, $V(i,j,k)$ and $W(i,j,k)$ but only $U(i+1,j+1,k+1)$, $V(i+1,j+1,k+1)$ and $W(i+1,j+1,k+1)$ • Therefore, in a small region around the point (i,j,k) the similarity function can be computed using a fixed deformation matrix.

Of course, it is only possible to compute $C(i,j,k)$ for a small number of values of (U,V,W) in any region. From the values of $C(i,j,k)$ at those points in the region we can get an analytic expression for $C(U,V,W)$ using the least square approximation. The simplest approximation - a function linear in U , V and W would result in constant force terms. When the values of (U,V,W) are close to the point of maximum similarity, the correlation function will have a peak. In

this case the direction of the force vector is changed as result of small changes in (U,V,W). This can cause 1 iterative process to oscillate around that point. Therefore, a quadratic approximation written as:

$$C \gg C_0 + C_u \cdot U + C_v \cdot V + C_w \cdot W + C_{uu} \cdot U^2 + C_{vv} \cdot V^2 + C_{ww} \cdot W^2 + 2C_{uv} \cdot U \cdot V + 2C_{vw} \cdot V \cdot W + 2C_{wu} \cdot W \cdot U.$$

is more appropriate*

When we substitute the approximations for 1 derivatives of U, V and W into the definitions of the force we get:

$$F_x(i,j,k) = [U(i+1,j,k) - U(i-1,j,k)] \cdot [C_{uu} \cdot U(i,j,k) + C_{uv} \cdot V(i,j,k) + C_{uw} \cdot W(i,j,k) + C_u / 2] + [V(i+1,j,k) - V(i-1,j,k)] \cdot [C_{uv} \cdot U(i,j,k) + C_{vv} \cdot V(i,j,k) + C_{vw} \cdot W(i,j,k) + C_v / 2] + [W(i+1,j,k) - W(i-1,j,k)] \cdot [C_{uw} \cdot U(i,j,k) + C_{vw} \cdot V(i,j,k) + C_{ww} \cdot W(i,j,k) + C_w / 2]$$

$$F_y(i,j,k) = [U(i,j+1,k) - U(i,j-1,k)] \cdot [C_{uu} \cdot U(i,j,k) + C_{uv} \cdot V(i,j,k) + C_{uw} \cdot W(i,j,k) + C_u / 2] + [V(i,j+1,k) - V(i,j-1,k)] \cdot [C_{uv} \cdot U(i,j,k) + C_{vv} \cdot V(i,j,k) + C_{vw} \cdot W(i,j,k) + C_v / 2] + [W(i,j+1,k) - W(i,j-1,k)] \cdot [C_{uw} \cdot U(i,j,k) + C_{vw} \cdot V(i,j,k) + C_{ww} \cdot W(i,j,k) + C_w / 2]$$

$$\begin{aligned}
F_z(i,j,k) = & [U(i,j,k+1) - U(i,j,k-1)] \cdot [C_{uu} \cdot U(i,j,k) \\
& + C_{uv} \cdot V(i,j,k) + C_{uw} \cdot W(i,j,k) + C_u / 2] \\
& + [V(i,j,k+1) - V(i,j,k-1)] \cdot [C_{uv} \cdot U(i,j,k) \\
& + C_{vv} \cdot V(i,j,k) + C_{vw} \cdot W(i,j,k) + C_v / 2] \\
& + [W(i,j,k+1) - W(i,j,k-1)] \cdot [C_{uw} \cdot U(i,j,k) \\
& + C_{vw} \cdot V(i,j,k) + C_{ww} \cdot W(i,j,k) + C_w / 2]
\end{aligned}$$

These terms introduce more coupling between the sets of equations for U, V and W and also make them nonlinear. The nonlinearity is not a serious problem, since the equations are solved by an iterative method which can handle it. It can only affect the convergence of the process. The problem caused by the presence of U(i,j,k), V(i,j,k) and W(i,j,k) in all the three equations for the point (i,j,k), is handled by solving in each step of the iteration the three equations for every point simultaneously.

The main problem with the quadratic approximation is caused by the fact that it does not always represent a function with a peak. That is, the surfaces of constant value of C are not always ellipsoids. While in principle at least this should not happen, in practice when the least square approximation is used, this can happen quite often. To handle these cases, we need the tools to detect them.

The position of the peak (or the extremum) of this function is given by:

$$\text{grad } C(U,V,W) = 0$$

The values of (U,V,W) at this point are the solution of the following system:

$$C_{uu} \cdot U + C_{uv} \cdot V + C_{uw} \cdot W = - C_u / 2$$

$$C_{uv} \cdot U + C_{vv} \cdot V + C_{vw} \cdot W = - C_v / 2$$

$$C_{uw} \cdot U + C_{vw} \cdot V + C_{ww} \cdot W = - C_w / 2$$

For these equations to have a solution the determinant of the coefficient matrix should not vanish. In real applications, the vanishing of the determinant is a very rare case, unless all the coefficients are zeros. Given the solution of these equations, it is possible to translate the origin of the coordinates to that location and to obtain an expression which does not contain the linear terms.

The mixed terms ($C_{uv} \cdot U \cdot V$, $C_{vw} \cdot V \cdot W$ and $C_{uw} \cdot U \cdot W$) can be eliminated by rotating the coordinate system. The coefficients of the quadratic terms that remain after this rotation are the eigenvalues of the above matrix. The quadratic expression for C represents an ellipsoid if all the eigenvalues of this matrix are positive. The necessary and sufficient conditions for this are:

$$C_{uu} + C_{vv} + C_{ww} > 0$$

$$C_{uu} \cdot C_{vv} + C_{vv} \cdot C_{ww} + C_{uu} \cdot C_{ww} - C_{uv}^2 - C_{vw}^2 - C_{uw}^2 > 0$$

$$C_{uu} \cdot C_{vv} \cdot C_{ww} + 2 \cdot C_{uv} \cdot C_{vw} \cdot C_{wu} \\ - C_{uu} \cdot C_{vw}^2 - C_{vv} \cdot C_{uw}^2 - C_{ww} \cdot C_{uv}^2 > 0$$

The last expression is the coefficients determinant. These conditions are obtained by considering the three invariants of the bilinear form and noting that their signs must be the same as in the diagonal representation.

When one or more of these eigenvalues is negative, the correlation function has a minimum along the direction of the corresponding eigenvector. This can happen in two cases. The first is when the region used to approximate C contains a real minimum of C . The other case is when the eigenvalue was small and became negative as a result of the approximation. A small eigenvalue often happens when the region contains an edge. In this case tangential components of the force relative to the edge can vanish.

The problem in using the quadratic expression for C is caused by negative eigenvalues. In this case the force becomes stronger as the distance from the extremum point becomes larger. This will cause unrealistically large changes in (U, V, W) in a single iteration step. To prevent this problem we need only set the value of the negative eigenvalue to zero. This process consists of finding the

rotation matrix that transforms the bilinear form into a diagonal one, setting the negative eigenvalues to zero and finally transforming the diagonal form into the original coordinates. This process consumes a relatively large amount of time, but in most cases it results in a faster convergence of the iteration process.

6.3 Three Dimensional Regression

Up to this point it was assumed that every point in the image is equally important. This assumption is justified in this work although for some points in the image it is easier to locate their corresponding point in the other image. These distinguished points usually have a sharper peak in the cross correlation function than the others. Therefore, the external forces that act on them are stronger than the external forces that act on the others when they are at the same distance from their 'true' placement.

There are however many applications where it is required to find the optimal mapping based on a set of a relatively small number of corresponding pairs of points in the two images, instead of the entire set of points in the image. These corresponding points could for example be located by special operators or could be entered as data

It is also possible in these cases that a different values of error estimates will be associated with each pair of data points. These error estimates specify the probability of finding the 'true' point (in one of the images) at a given distance from the data point. When the probability distribution is symmetric its shape is usually described by the Gaussian:

$$\text{EXP} (- r^2 / \sigma^2)$$

where r is the distance from the data point.

The assumption about the symmetry of the probability distribution is often made because of the difficulties in handling the more general case of asymmetric distribution given by:

$$\text{Pr}(U,V,W) = \text{EXP} (- E(U - U', V - V', W - W'))$$

where E is a second degree polynomial representing ellipsoids centered at (U', V', W') (*). This type of registration problem is very common in many areas. Tobler called this problem in the two dimensional case by the name of bidimensional regression [TOBLER 1977], and gave an ad hoc method for obtaining a smooth solution without considering the errors.

(*The conditions that the surfaces obtained from $E = \text{const}$ are three dimensional ellipsoids are given in section 5.2.

Our method can be used to obtain an optimal mapping with or without these error estimates. Before delving into the details let us first discuss the cost function to be minimized by the optimal mapping. We will first assume that no error estimates are present, we will also leave out the deformation part of the mapping and assume that the mapping is represented by three polynomials P_u , P_v and P_w . let us denote by (X_m, Y_m, Z_m) the coordinates of the point in one image and by (U_m, V_m, W_m) the coordinates of its corresponding point in the other image. Using the least squares approximation means the minimization of the following error function:

$$\text{ERROR}^2 = \sum_{m=1}^N \left\{ [U_m - P_u(X_m, Y_m, Z_m)]^2 + [V_m - P_v(X_m, Y_m, Z_m)]^2 + [W_m - P_w(X_m, Y_m, Z_m)]^2 \right\}$$

When error estimates are available and the probability distribution is symmetric each term in the sum is multiplied by a weight which is the width of the corresponding Gaussian. When the distribution is asymmetric but the error function becomes

$$\text{ERROR}^2 = \sum_{m=1}^N E(U_m - U_m^* V_m - V_m^* W_m - W_m^*)$$

m*1

where E is the same second degree polynomial that appears in

the exponent of the probability distribution.

The cost function that is minimized in our method is

$$\text{COST} = \text{DEFORMATION} + \text{ERROR}^2$$

As in the matching, we associate the second term with the potential function from which the external forces are derived. The solution of this problem is obtained by solving the three partial differential equations without the need to give an explicit expression for the deformation part. The solution of these equations is again obtained by the finite difference approximation using a rectangular grid of points.

If the point (X_m, Y_m, Z_m) falls on a grid point then we can easily get the components of the external forces at that point. More often however, the data points will not fall on the grid points but in a cell whose vertices are the grid points. In such a case some sort of interpolation is required. Let (i, j, k) be a grid point such that:

$$X(i-1, j, k) < X_m < X(i+1, j, k)$$

$$Y(i, j-1, k) < Y_m < Y(i, j+1, k)$$

$$Z(i, j, k-1) < Z_m < Z(i, j, k+1)$$

we will use the following notation:

$$\Delta X_m = [x_m - x(i, j, k)]$$

$$\Delta Y_m = [y_m - y(i, j, k)]$$

$$\Delta Z_m = [z_m - z(i, j, k)]$$

$$U_x(i, j, k) = [U(i+1, j, k) - U(i-1, j, k)] / 2$$

$$U_y(i, j, k) = [U(i, j+1, k) - U(i, j-1, k)] / 2$$

$$U_z(i, j, k) = [U(i, j, k+1) - U(i, j, k-1)] / 2$$

and similar expressions for v_x, v_y, v_z and w_x, w_y, w_z .

Using linear approximation we can write:

$$U(x_m, y_m, z_m) = U(i, j, k) + \Delta X U_x(i, j, k) + \Delta Y U_y(i, j, k) + \Delta Z U_z(i, j, k)$$

$$V(x_m, y_m, z_m) = V(i, j, k) + \Delta X V_x(i, j, k) + \Delta Y V_y(i, j, k) + \Delta Z V_z(i, j, k)$$

$$W(x_m, y_m, z_m) = W(i, j, k) + \Delta X W_x(i, j, k) + \Delta Y W_y(i, j, k) + \Delta Z W_z(i, j, k)$$

The components of the external force at (i, j, k) are given by:

$$F_x(i, j, k) = \frac{\partial E}{\partial U_m} U_x(i, j, k) + \frac{\partial E}{\partial V_m} V_x(i, j, k) + \frac{\partial E}{\partial W_m} W_x(i, j, k)$$

$$F_y(i, j, k) = \frac{\partial E}{\partial U_m} U_y(i, j, k) + \frac{\partial E}{\partial V_m} V_y(i, j, k) + \frac{\partial E}{\partial W_m} W_y(i, j, k)$$

$$F_z(i, j, k) = \frac{\partial E}{\partial U_m} U_z(i, j, k) + \frac{\partial E}{\partial V_m} V_z(i, j, k) + \frac{\partial E}{\partial W_m} W_z(i, j, k)$$

where the expressions for $U(X_m, Y_m, Z_m)$, $V(X_m, Y_m, Z_m)$ and $W(X_m, Y_m, Z_m)$ are substituted after the differentiation for U_m , V_m and W_m respectively. These expressions are linear in $U(i, j, k)$, $V(i, j, k)$ and $W(i, j, k)$ but are quadratic in the others. When there are several data points within the grid cells that have the point (i, j, k) as one of their corners the total force is simply the sum of their contributions. Grid points that are not corners of cells with data points have no external force to act on them. Their positions will be determined by the internal forces only.

To solve the equations obtained in this way we use an iterative method. Again, the three equations for each grid point are solved simultaneously to increase the convergence rate of the method. Unlike the case where the 'true' position of each grid point is unknown and the correlation function has to be computed again in each step using a lengthy process, here we can store the expressions for the forces and just recompute them each time using the new values of the unknowns.

There is only one more point that should be mentioned here. If the Gauss - Siedel iterative method is used, it is possible that the placement of one point will be moved so far from its previous placement, as to change the sign of the partial derivatives. To avoid any problems from this Jaccobi's method should be used.

0.4 Boundary Conditions

The Dirichlet boundary conditions specifying the U , V and W at the boundary grid points are the easiest to use. They are also the most restrictive, i.e. they have the largest effects on the solution. These values can be obtained from the global mapping. Since the global mapping is only a linear approximation to the optimal mapping, it is not necessarily a good one for every point on the boundary. As a result, one can expect to get larger errors near the boundary which can affect the entire registration.

There is, however, a way to reduce their effects on the solution in the region of interest. One simply has to make the grid larger than the image, so that the boundary grid points whose values are fixed will be far away from the grid points where the solution is important. The price for this solution is paid by the additional computations required by the extra points.

The Neumann boundary conditions of specifying the derivatives of U , V and W at the boundary grid points have the advantage in our application. These conditions can also be obtained from the global mapping. Thus, if the global mapping is given by:

$$\begin{pmatrix} U \\ V \\ W \end{pmatrix} = \begin{pmatrix} G_{10} \\ G_{20} \\ G_{30} \end{pmatrix} + \begin{pmatrix} G_{11} & G_{12} & G_{13} \\ G_{21} & G_{22} & G_{23} \\ G_{31} & G_{32} & G_{33} \end{pmatrix} \cdot \begin{pmatrix} X \\ Y \\ W \end{pmatrix}$$

then

$$\frac{\partial U}{\partial X} = G_{11}, \quad \frac{\partial U}{\partial Y} = G_{12}, \quad \dots, \quad \frac{\partial W}{\partial Z} = G_{33}$$

In our implementation, the boundary points of the grid lie on the faces of a rectangular box. Assuming that one of these faces contains the points $(0, j, k)$, the X-derivative of U on this face can be approximated by:

$$\frac{\partial U(0, j, k)}{\partial X} \doteq [U(1, j, k) - U(0, j, k)] \doteq G_{11}$$

Similar expressions can be written for the X-derivatives of V and W. It seems like we have obtained one equation for each additional unknown (in this example: $U(0, j, k)$), but in fact, the system of all the equations does not necessarily have a unique solution (*). If however the forces are functions of (U, V, W) as is the case here, a unique solution will result*

Since only one boundary condition is sufficient for obtaining the value of $U(0, j, k)$ it is clear that we cannot use the other conditions. This can result in large

(*) This a general property of the Laplace equations.

displacements of the boundary points. To minimize this problem we change the above condition into:

$$U(0,j,k) = [U(1,j,k) - G_{11}] / 3 + [U(0,j+1,k) + U(0,j-1,k) + U(0,j,k+1) + U(0,j,k-1)] / 6$$

which makes $U(0,j,k)$ be the average of the above requirement and the following two:

$$\frac{\partial^2 U}{\partial Y^2} = 0, \quad \frac{\partial^2 U}{\partial Z^2} = 0.$$

At the edges of the boundary box, for example, at the points $(0,0,k)$ the conditions used in our implementation are:

$$U(0,0,k) = [U(1,0,k) - G_u + U(0,1,k) - G_{12}] / 3 + [U(0,0,k+1) + U(0,0,k-1)] / 6$$

Finally, the conditions at the vertex $(0,0,0)$ of the boundary box is:

$$U(0,0,0) = [tU(1,0,0) - G_u + U(0,1,0) - G_{12} + U(0,0,1) - G_{13}] / 3$$

Similar conditions are used for V and W and for the other boundary points.

CHAPTER 7 IMPLEMENTATION

7.1 General Considerations

The registration method described in this work is a general one and can be used to register any pair of similar images. The implementation of this method was however tailored to our particular application. While in other application areas, each time a new pair of images is registered, in our case, one image of the pair is always the same and only the other one is new. The fixed image is the atlas and the new one is the CT image. It is possible to take advantage of this fact to reduce the amount of computation required by the process.

The most time consuming action of our registration process is the computation of the projections of the image function on the orthogonal base. Since the atlas image is always the same, its projections need to be computed only once and they can be saved on a disc or a magnetic tape. Because of this, the process was designed to use projections on more points in the atlas image than points in the CT image.

While it is possible to store the projections of a many regions as one may want, there is a price for it. Because the size of the computer memory is limited, it is not usually possible to keep the entire set of projections in memory and therefore an overlay structure is required. Even when large computers are used the memory is paged, and only some of the pages are kept in the main memory. Since the projections are used again and again, they have to be read from the paging device many times. This operation sets a limit to the number of regions that could be used.

The number of regions in the CT image is the same as the number of (internal) grid points. This number is determined by the image size and the mesh size. The resolution of the mapping depends on the mesh size since the mapping of points other than the grid points can only be found by interpolation. A lower limit on the effective mesh size is imposed by the resolution of the image through the size of the correlation regions.

In order to take into account the deformation of the image in the similarity measurements, the projections of the image function on the orthogonal base should be modified (or recomputed) according to the deformation. The deformation within a region is approximated in our method by the linear deformation matrix and the elements of this matrix are approximated by finite differences. For each grid point

(if $j > k$) these approximations contain the mapping values at the $(i \pm 1, j \pm 1, k \pm 1)$ grid points. Therefore, the radius of the spherical region in which the linear approximation is valid is close to the distance between two grid points. This radius should on the other hand be greater than the resolution of the CT image or else only the projection 01 the constant function (H_0) will not be identically zero. The thickness of the CT slices used to test our method was 8mm. Therefore, we have used a grid with an 8mm mesh size and regions with a radius of 12mm.

The number of useful projections for each region depends on the size of the region and on the amount of detail (texture) within a region. The magnitudes of the projections decrease with the number of zero crossings (frequency) of the corresponding base functions. The relative contribution of the different projections to the cross correlation function is proportional to the square of their values. To limit the error in this function to about one percent, only those projections whose magnitude is greater than one tenth of the largest projections are required. Being limited by the size of the computer memory we have mostly used only the first four projections from the set given in section 5.3. Increasing the number of projections did not produce significant effects on the resulting mapping.

7.2 The Global Mapping

The first step in the registration process is to find the global mapping*. This is done by placing a grid on the CT image and computing the projections for each grid point. We assume that the two images are already grossly registered so that both the scales and the orientations- of the two are similar. If this is not the case, then the grid placed on the CT image is scaled and rotated according to our best knowledge•

Based on our estimate of the similarity between the two images we can set a limit to the displacement between the two. This limit is used to define the size of the search area in the atlas. For each point in the grid a search is conducted to find the most similar point within the search area in the atlas. The similarity is measured by the cross correlation of the corresponding projections. A weight is assigned to each internal grid point by:

$$w(i,j,k) \ll \text{EXP}(-d^2(i,j,k))$$

where $d(i,j,k)$ is the distance between the best position of the point (i,j,k) and the average positions of the best points obtained for its six neighbors.

From these weighted positions a global mapping is computed using the least squares approximation. If the deformation matrix of this mapping is close to the unit matrix, the global mapping is accepted and the elastic part of the registration process can begin. Otherwise, a new grid based on this mapping is put on the CT image and the projections are computed using the new grid as the coordinate system. Using the new projections the search process is repeated and a new deformation matrix is computed again.

The iterative process of finding the global mapping is repeated until there is (almost) no difference between two successive mappings. Recomputing the new projection is required only if the global mapping contains large deformations. If the global mapping is mostly a rotation and translation the new projections can be easily obtained from the old ones using the procedure described in section 5.4.

It should be noted that the reliability of this process depends to a great extent on the images themselves. If the images contain sufficient details so that many points can be uniquely identified during the search process, the resulting mapping is quite good. In other cases the particular point selected as the most similar one to any grid point depends on the order in which the search is conducted. A useful

strategy is to select from the set of equally similar points the one which requires the smallest displacement.

7.3 Local Matching

After the global mapping is either computed or estimated, the main step of the registration process can start. The main problem in this step is to find in each iteration cycle and for each grid point the components of the external force. To find these components we need in this method to find the quadratic approximation to the similarity (cross correlation) function.

Since it is possible to evaluate this function only for those points for which the projections are available, we have to use the least squares approximation method. In order that this approximation will be useful it should represent a function that has a maximum point near the 'true' placement of the grid point. The assumption that the cross correlation function can be approximated by a quadratic expression with negative eigenvalues (*) is usually valid only in a small region around the maximum. When values at points which are far from the peak are used,

(*) The eigenvalues are related to the second derivatives of the function which are negative at the point where the function attains its maximum.

the quadratic approximation obtained from the least squares method can represent a variety of functions with undesired properties for our process. If the peak of the function is too narrow compared with the mesh size of the atlas grid, this situation will occur frequently. Making the mesh size smaller will force us to conduct a lengthy search for the peak. Another related problem is that for many points instead of a peak the function has a plateau, and therefore there is no point in searching it.

A method was therefore required to determine whether the program should search for a peak and if so a way to reduce the search time. The cross correlation function has a peak only if there is one in the autocorrelation function. This function could have a peak if the sum of the magnitudes of projections other than I_0 (the projection on the constant function) is greater than zero. The larger this sum is compared with I_0 , the more narrow is the peak. This property can be used to sharpen the peak by omitting I_0 from the cross correlation or to make it wider by assigning a larger weight to I_0 and smaller ones to the others. In the extreme case we would like to use only the I_0 term, but in this case we have to compute the similarity by the absolute value of the difference between the I_0 's at the two points.

When the peak is wide, a coarse grid can be used to approximate the similarity function. If the quadratic approximation obtained in this way has negative eigenvalues, we can find the position of its maximum. If it does not have a peak, no external forces will be applied to that grid point. Using a fine grid around this point and a small weight for I_0 , the program recomputes a new quadratic approximation for the similarity function. Of course, the same atlas grid is used for these two approximations, but in the former case the program uses every Nth point in it.

When the second approximation is computed, the signs of its eigenvalues are tested. A positive eigenvalue can cause the iteration process to diverge as the force in the direction of the corresponding eigenvector pulls the point away. This force becomes larger and larger as the solution runs away from the right place. Therefore a positive eigenvalue should be zeroed if the approximation is to be used. Since a transformation that can cause round off errors is applied to the coefficients of the similarity function, the program changes the positive or zero eigenvalues to a small negative value.

The number of points used by the least squares method is $4^3 = 64$ in the three dimensional implementation and $4^2 = 16$ in the two dimensional case. It is easier to compute the approximation and also the errors are smaller if

a coordinate system in which the points are symmetrically distributed around the origin. The change to another coordinate system is trivial.

It is clear that this process takes a considerable amount of time and therefore its use should be economized. Since in each iteration step the placement of each grid point can only be changed by a small amount as its neighboring points are holding it back, the same approximation can be used during several iterations. When the iterative process converges, new approximations are computed. These are used to improve the mapping by further iterations.

To obtain the improved approximations the cross correlation function is computed using modified values for the projections. These modifications are done by the method described in section 4.4. using the approximated mapping to find the rotation vector. If the amount of pure deformation is too large they should be recomputed by changing the region shape to the appropriate ellipsoid and deforming the orthogonal base functions. In our implementation recomputations were not required. By avoiding this process we could divide the process into several tasks that could fit the available memory.

The equilibrium position for each point in each iteration is first computed in the absence of external forces. The similarity function is transformed to a coordinates system in which this point is the origin. The components of the external force are obtained in this system. The three equilibrium equations rewritten for the displacements from this origin are solved simultaneously, and the solutions are saved until the equations for all the other internal grid points are also solved (Jaccobi's iteration method). This prevents the possibility of obtaining the wrong sign for the derivatives of the mapping which could happen if a new placement for a grid point cross over an old placement of its neighbor. At the end of the iteration cycle the new values are also used to compute the mapping of the boundary points.

7.4 The Elastic Constants

Selecting the values of the elastic constants is also a point which deserves some discussion. The partial differential equations and as a result the finite difference equations can be simplified by setting $\lambda + \mu = 0$. When the two images to be matched contain many details, the solution is influenced mainly by the external forces. In such cases this simplification does not have undesired effects. When

the images contain large areas lacking details the process behaves in the two dimensional case differently than in the three dimensional case.

The reasons for this can be seen by considering Hooke's law in a coordinate system in which the strain matrix is diagonal. In the two dimensional case this law is written as:

$$S_1 = \lambda \cdot (e_1 + e_2) + 2 \cdot \mu \cdot e_1$$

Setting $\lambda + \mu = 0$ we get:

$$S_1 = \mu \cdot (e_1 - e_2) \quad \text{and} \quad S_2 = \mu \cdot (e_2 - e_1)$$

Thus in the absence of any stress deformations are still possible provided that they have the same sign and magnitude. This phenomenon is of course very undesirable because it can cause spontaneous shrinking or expansion in regions lacking details.

Setting $\lambda + \mu = 0$ in the three dimensional case, we will get:

$$S_1 = \mu \cdot (e_1 - e_2 - e_3)$$

$$S_2 = \mu \cdot (-e_1 + e_2 - e_3)$$

$$S_3 = \mu \cdot (-e_1 - e_2 + e_3)$$

and the only solution when

$$S_1 = S_2 = S_3 = 0$$

is:

$$e_1 = e_2 = e_3 = 0.$$

The behavior of an object with this property is certainly strange as can be seen by setting

$$S_1 = S, S_2 = S_3 = 0.$$

In response to a force pushing in one direction the object will shrink in the other two directions.

Setting $\lambda = 0$ is the next simple choice as only one constant is left. We have found that in two dimensional case a negative value for λ ($-\mu < \lambda < 0$) can sometimes produce better looking results. The three dimensional examples that were tested contained enough details to allow us to use $\lambda = -\mu$ without problems.

A suitable value for μ depends also on the amount of detail in the images. When the amount of detail is large the peak of the cross correlation function is narrow and the forces are strong. A small value for μ in this case could cause problems if the program finds a peak in the wrong place. This is more likely to happen during the first few iterations. Thus the best way is to start with a large value for μ and to decrease it as the process converges.

CHAPTER 8 RESULTS

The optimal registration method developed in this work was tested by registering (matching) several pairs of two and three dimensional images* The results of these tests are presented in this chapter.

In each example, one of the images is called the reference image and the other is called the test image. The reference images contain several objects each of them is a connected set of pixels having the same density (gray) value. In the two dimensional cases an object is defined by a set of points on its outer boundary. The boundary of the object is obtained by joining these points with straight lines. The object consists of all the pixels inside this boundary•

In one of the three dimensional examples, the objects were defined as set of pixels whose coordinates satisfy certain inequalities. In the other three dimensional example the reference image (the atlas) was constructed by stacking a set of two dimensional images (slices).

In the tests, a grid with a mesh size of eight pixels was placed on the test images and the optimal placement of these grid points were computed by the registration process. The placement of points other than the grid points were computed using linear interpolation. The inverse mapping, that is the point in the test image which corresponds to a given point in the reference image, is obtained by locating the point in the test image whose mapping point is the nearest one to the point in the reference image.

In the two dimensional examples the inverse mapping was used to deform a rectangular grid representing the reference image (the reference grid) which was then superimposed on the test image. The inverse mapping was also used to find the mapping of the boundary points defining the objects. The superimposed boundaries were obtained by drawing lines between consecutive boundary points.

A different method of presentation was used in the three dimensional examples. Pictures of the three dimensional structures from several view points were generated by making some of the objects opaque and the others transparent. The reconstructed images were obtained from the reference image using the mapping computed in the registration processes, and the results were displayed in the same way as the reference images.

synthetic images shown in figure 1, was used. The right hand side of figure 1a contains the reference image and the left hand side contains the test image that was registered with it. Figure 1b contains the result of the registration process. The left hand side of figure 1b shows the deformed reference grid superimposed on the test image. The right hand side of figure 1b shows the deformed boundaries of the objects in the reference image superimposed on the test image.

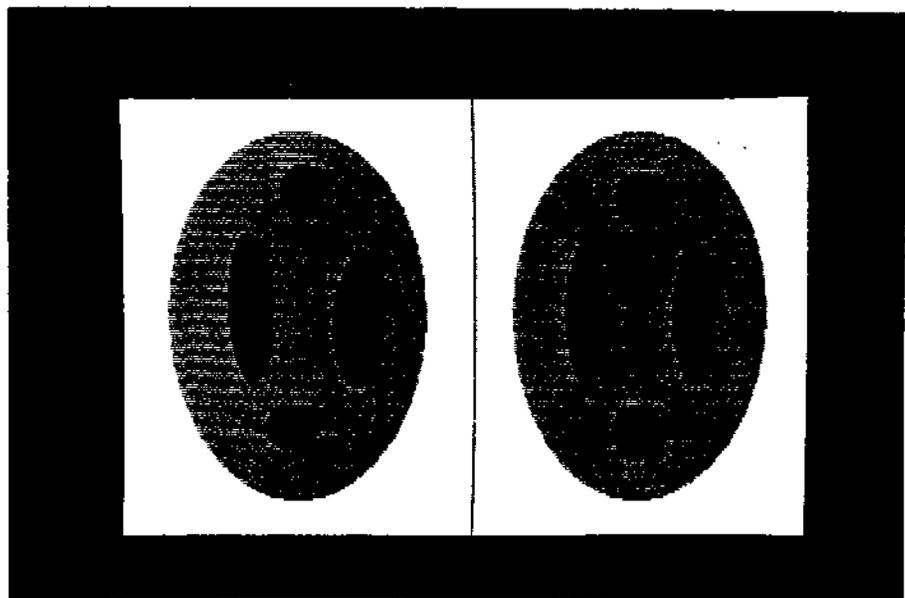


Figure 1a. Synthetic images for registration.

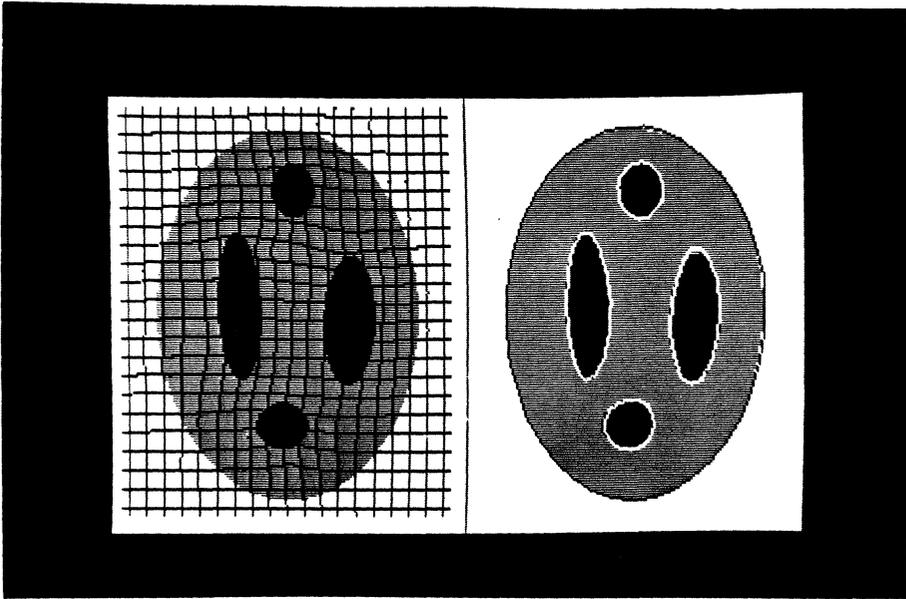


Figure 1b. Results of registration.

The same pair of synthetic image used in the first example was also used in the second example, but this time, noise was added to the test image. The gray values of the white background, the large gray ellipse and the four small black objects are: 200, 125 and 50 respectively. The noise level added to each pixel of the test image was a randomly generated number in the range $[-50, 50]$. Figure 2a shows the noisy test image (left hand side) and the reference image. Figure 2b contains the results of matching the noisy test image with the reference image. There is hardly any difference between these results and those shown in figure 1b.

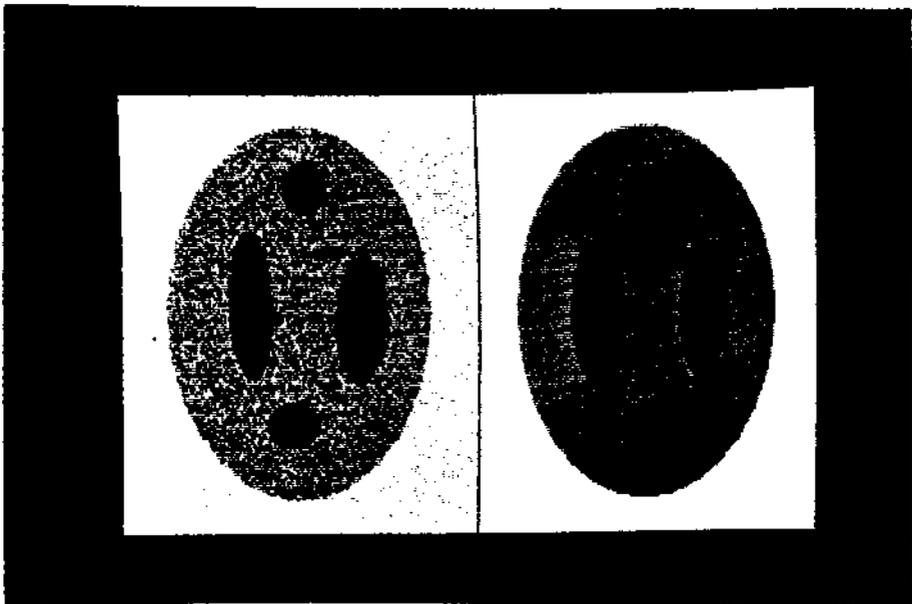


Figure 2a. Synthetic images with noise.

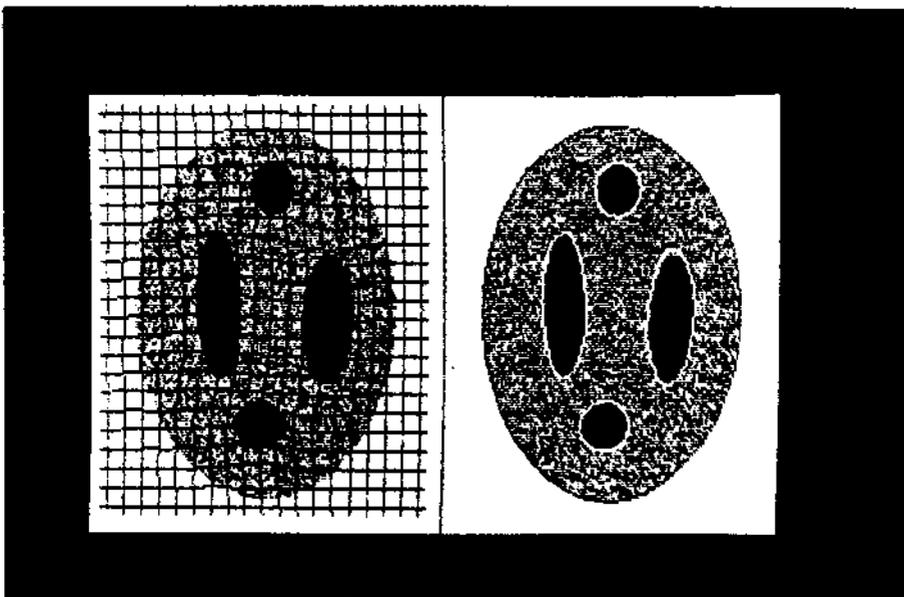


Figure 2b. Results of registration.

A different pair of synthetic images was used in the next example shown in figure 3a, where the test image is on the left side and the reference image is on the right side. The results of the registration process are shown in figure 3b. A large amount of grid deformation can be seen in the lower right region of the gray ellipse. As a result of this, there is some disparity between the superimposed boundary and the edge of the ellipse in this region.

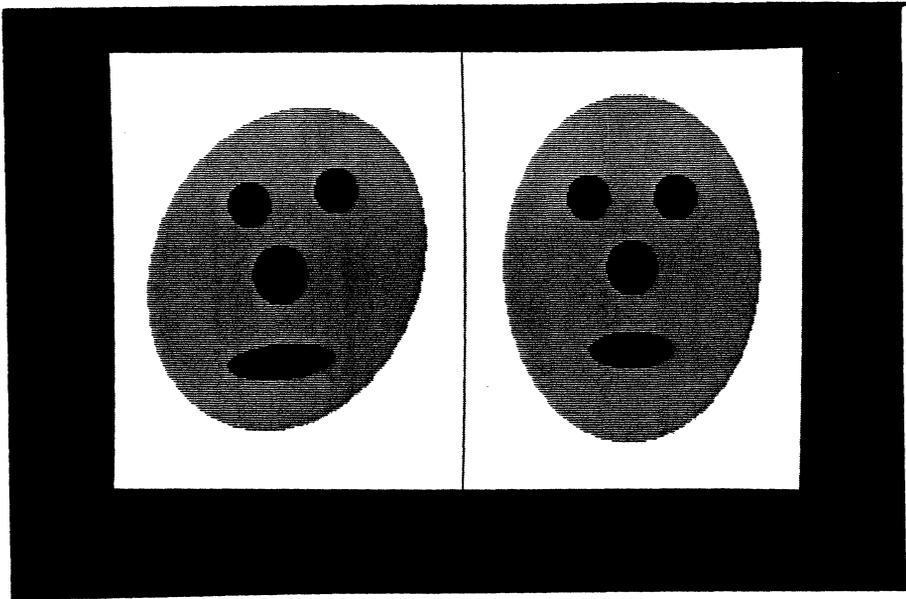


Figure 3a. Synthetic images for registration.

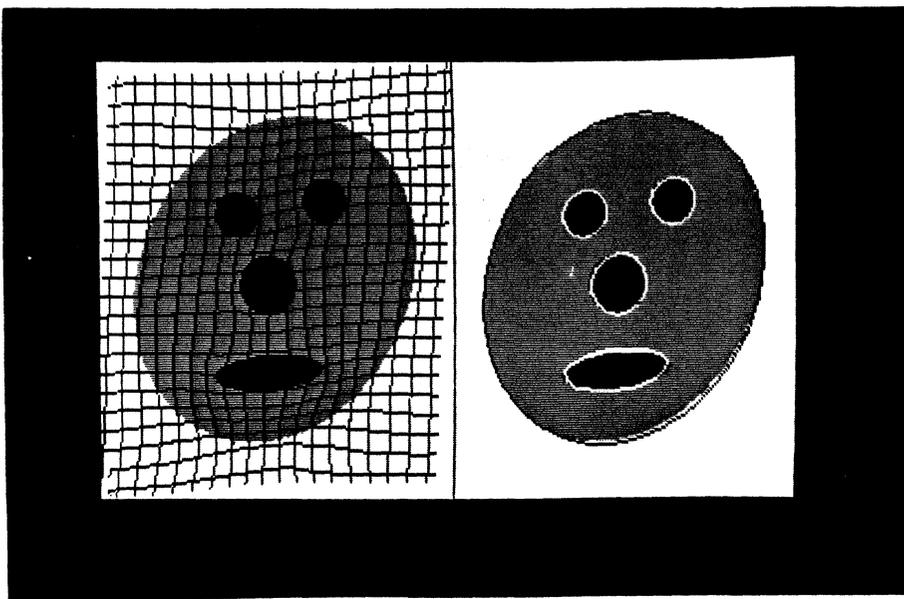


Figure 3b. Results of registration.

Figure 4a shows the same two images from the previous example except that a noise was added to the test image. Figure 4b shows the results of registering this pair of images, and like the first example, they are almost the same as those obtained for the clean test image.

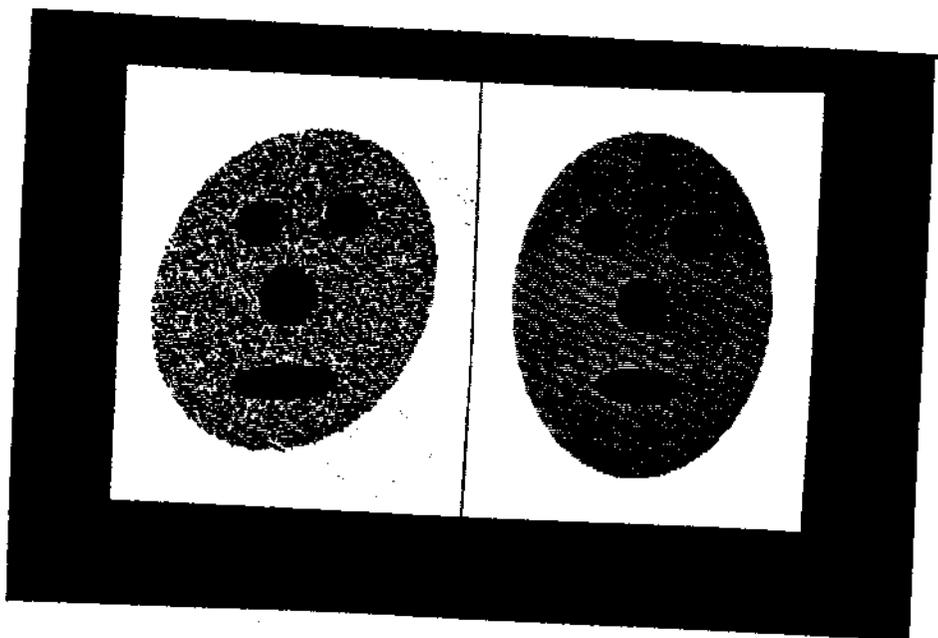


Figure 4a. Synthetic images with noise.

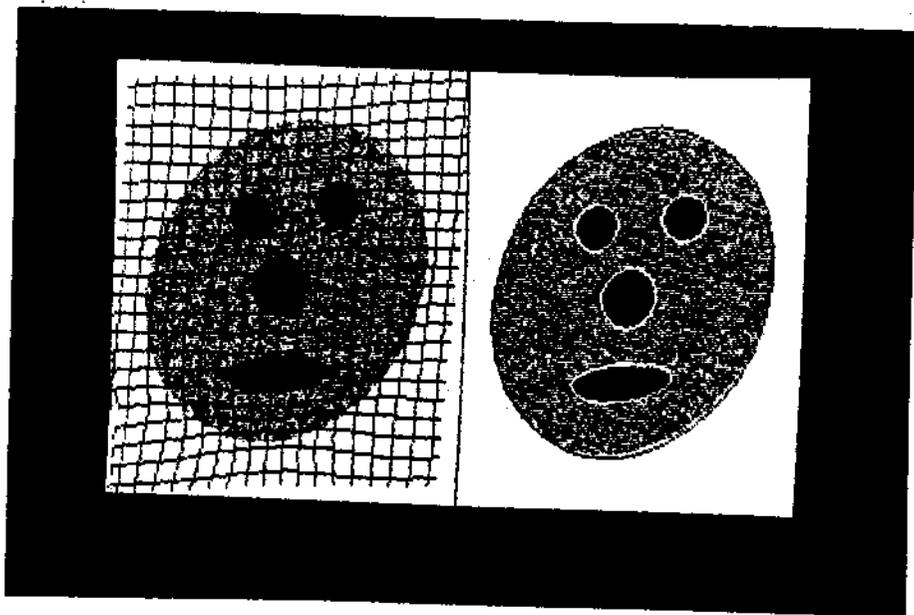


Figure 4b* Results of registration.

Figures 5a, 6a, 7a and 8a contain four CT images of a human brain (the left hand side of each figure) together with their corresponding atlas (reference) images (the right hand side). The atlas images contain the ventricles, the caudate, the thalamus and the putamen. The rest of the brain was taken as one object with a single density value. The regions around the brain, both in the CT and in the atlas images have the density value of a bone. Figures 5b, 6b, 7b and 8b show the results of registering these pairs. In most cases the boundaries of the anatomic objects are properly placed. Some misregistrations occurred in regions where the atlas and the CT images have different density values (for example, at the boundary of the brain near the central sulcus and the lateral sulcus where the CT images contain dark regions).

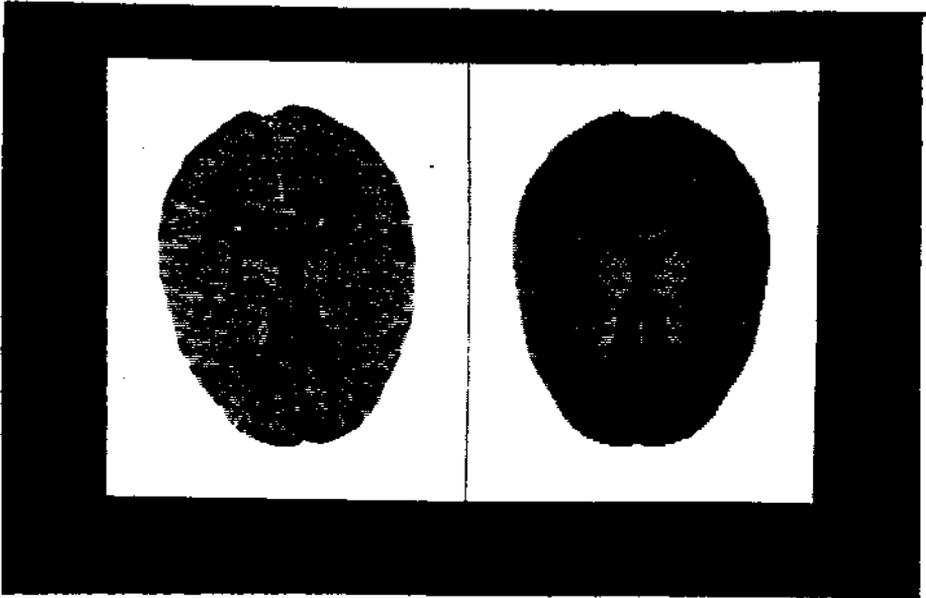


Figure 6a. A CT slice with its atlas image.

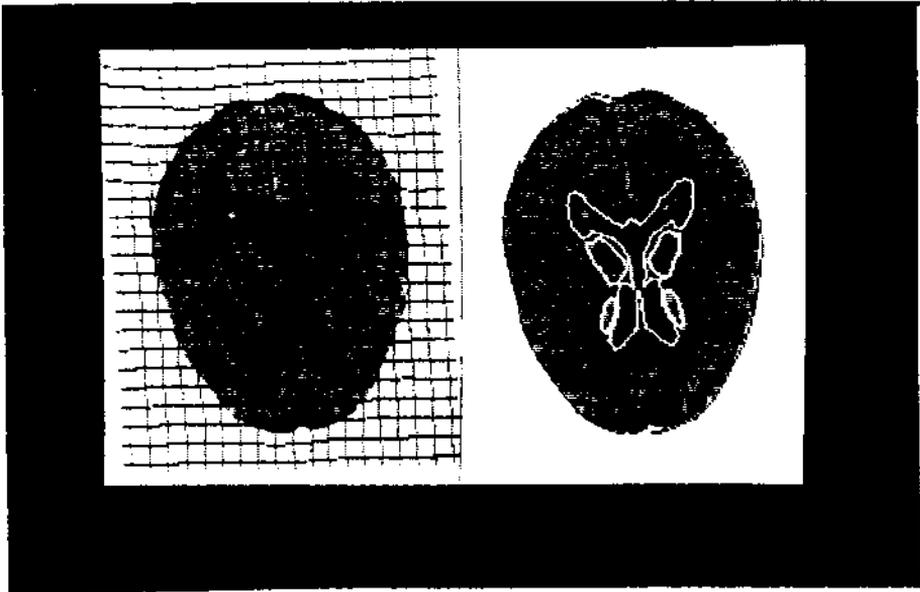


Figure 6b. Results of registration.

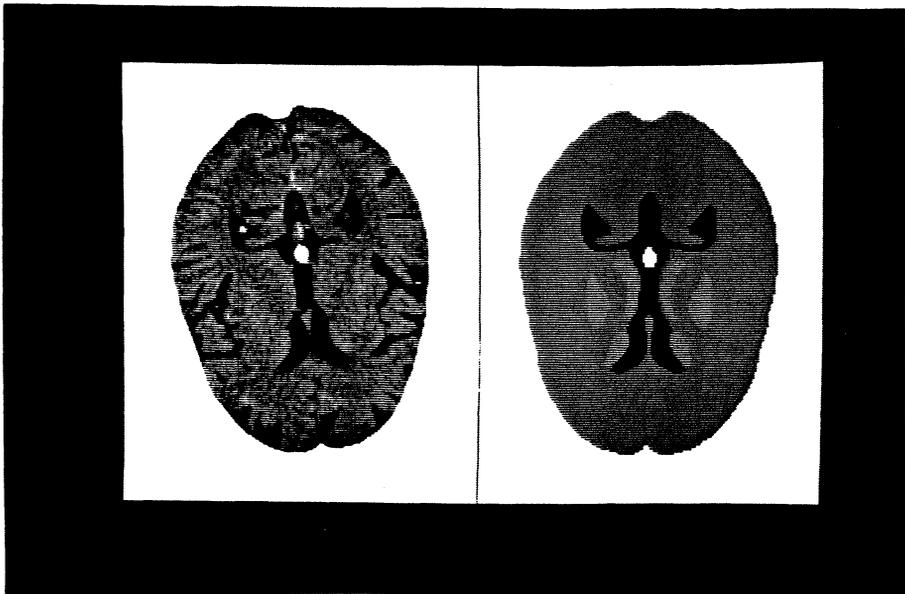


Figure 7a. A CT slice with its atlas image.

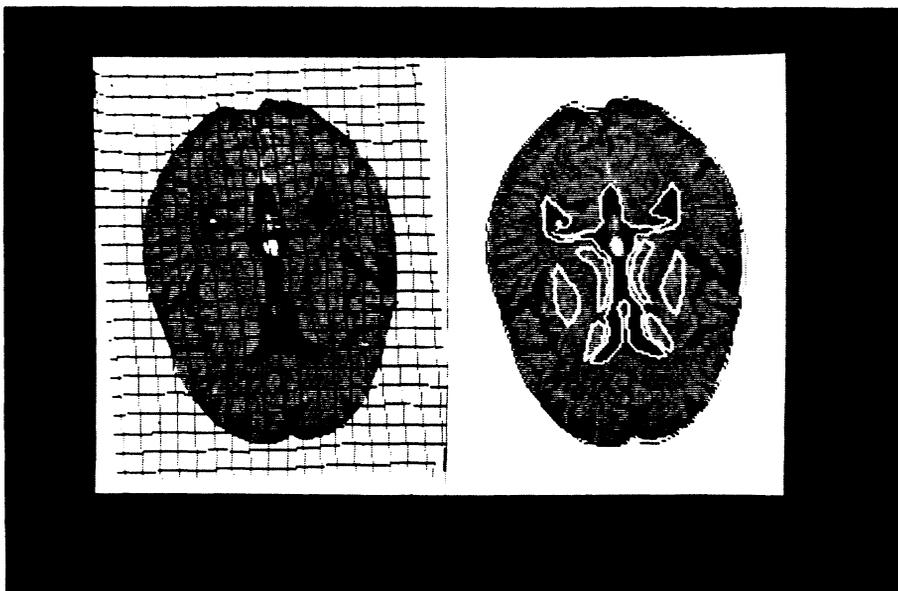


Figure 7b. Results of registration.

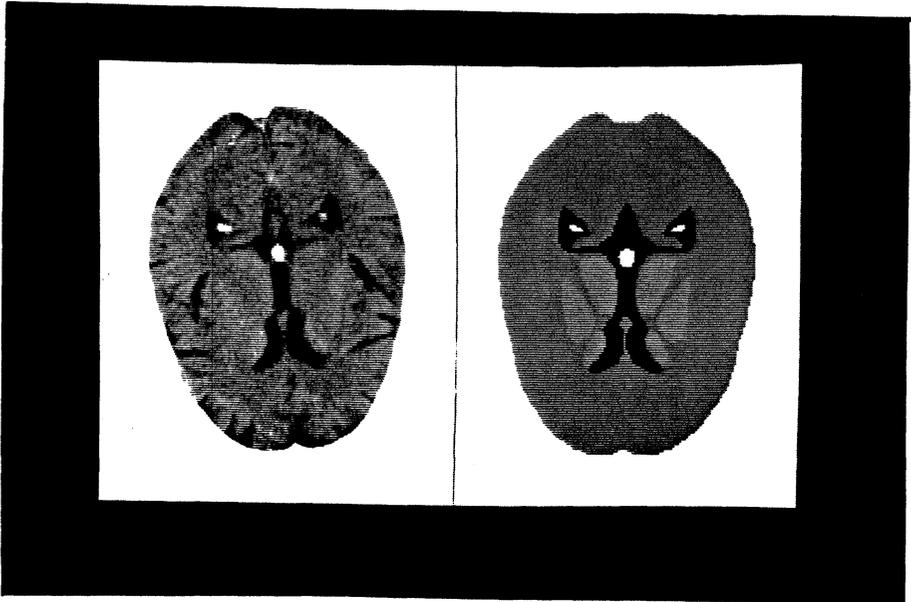


Figure 8a. A CT slice with its atlas image.

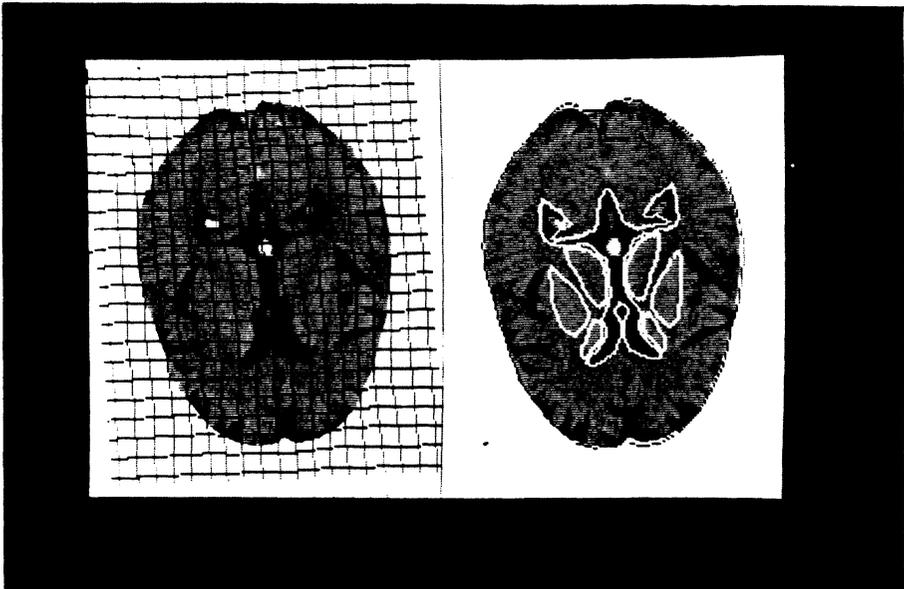


Figure 8b. Results of registration.

The next two examples are of the three dimensional registration. In the first one synthetic images were used both for the reference image and for the test image. Figures 9a-9c show three views of the reference object and the test object. Each object contains a barbell inside a torous within an ellipsoid. The test image differs from the reference image by the lengths of the ellipsoid axes, by the shape and orientation of the torous, and by a rigid transformation (a translation and a rotation) of the barbell.

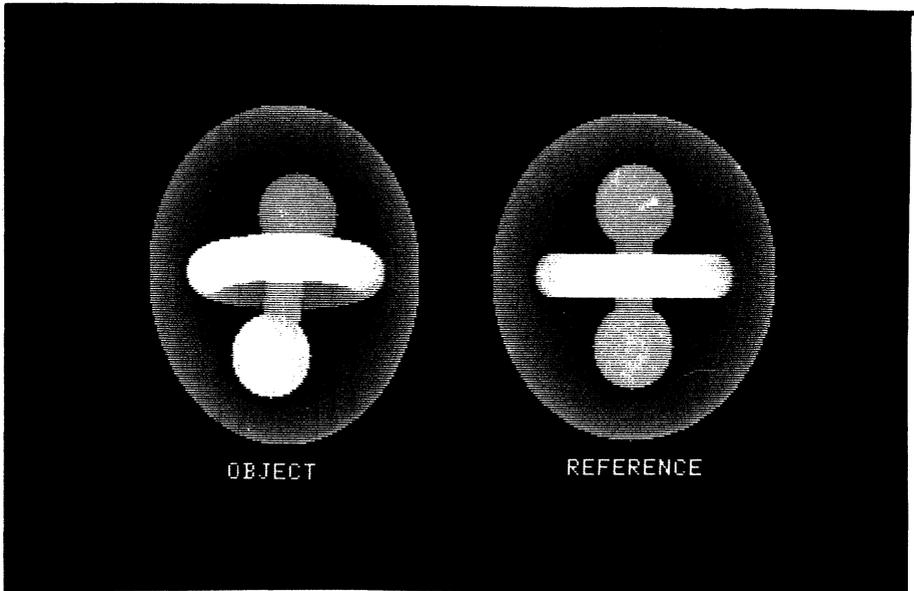


Figure 9a. Three dimensional objects.

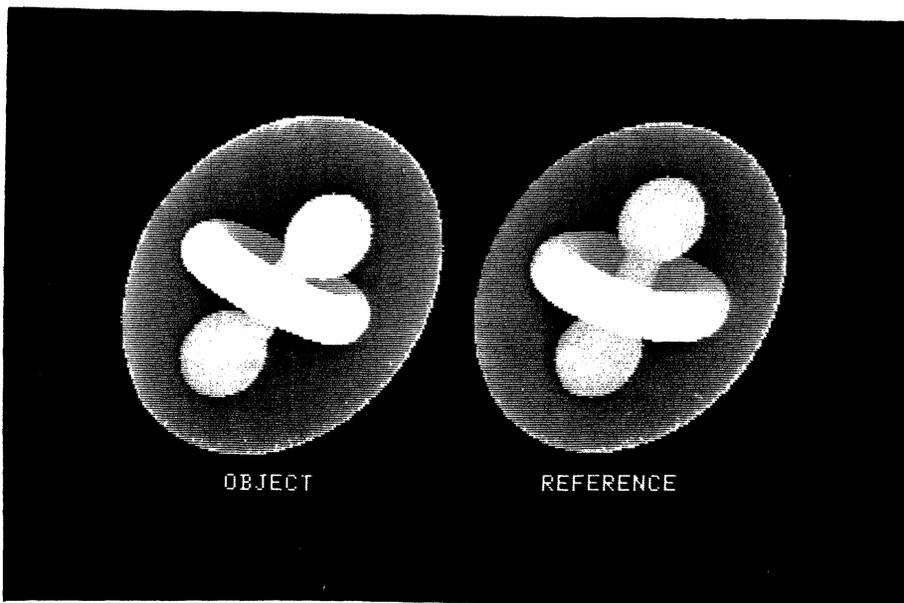


Figure 9b. Three dimensional objects.

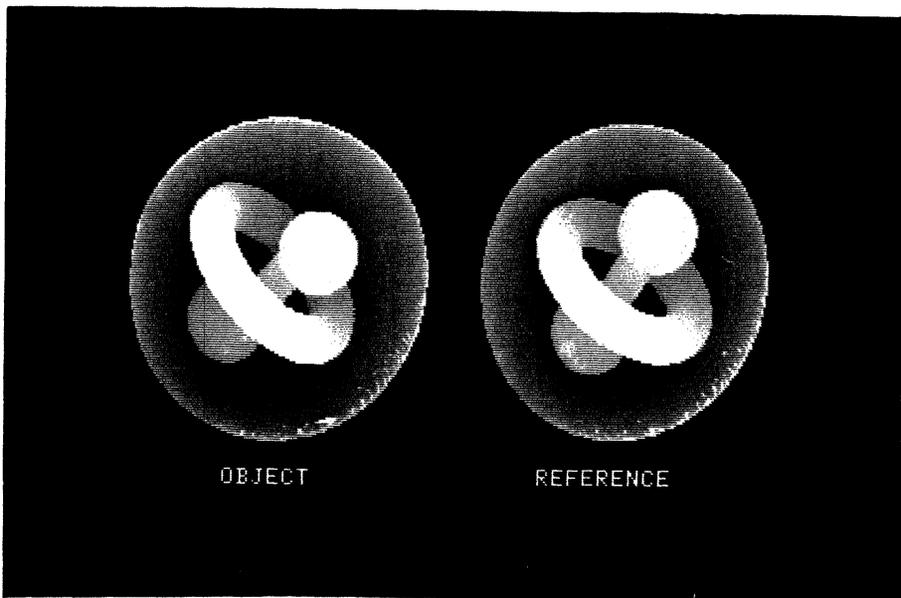


Figure 9c. Three dimensional objects.

Figures 10a-10c show the six slices of the test object (image) that contain the torous. These simulated CT images were created by averaging the densities of slices whose thickness were 8 pixels. An image made of 20 such slices stacked on top of each other was registered with the reference object. Figures 11a-11c show three views of the deformed reference object (the reconstructed images) next to the original test object. Figures 12a-12c show three views of the reconstructed object next to the reference object.

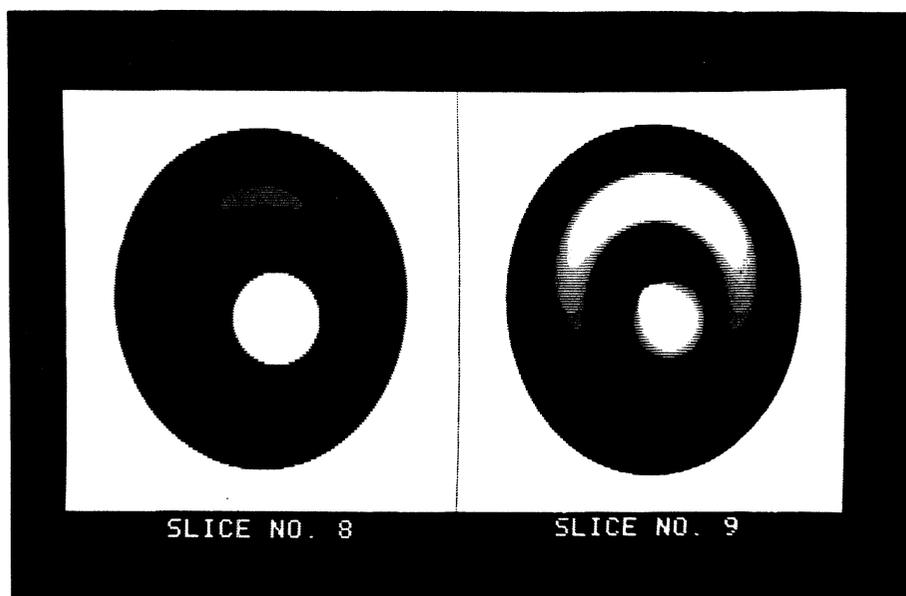


Figure 10a. Simulated CT slices.

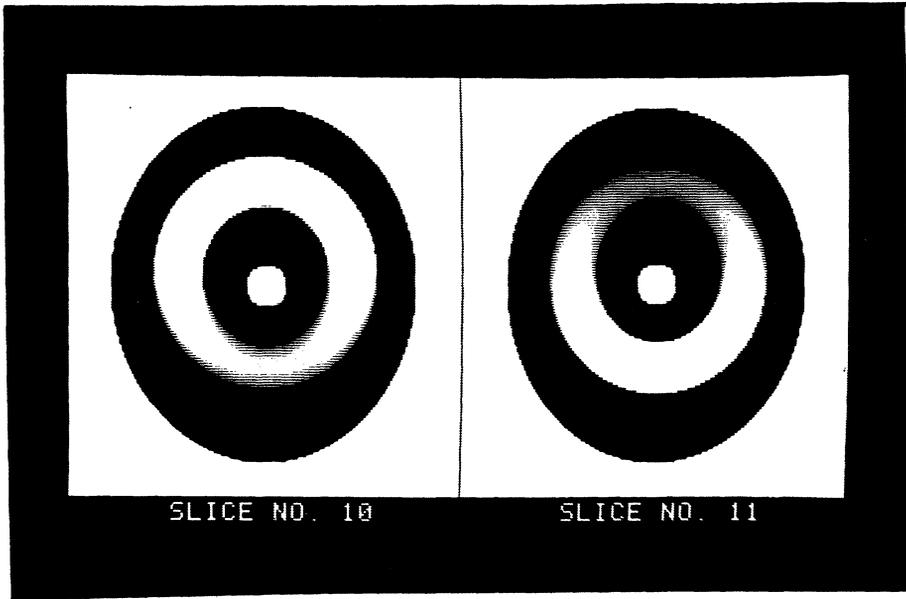


Figure 10b. Simulated CT slices.

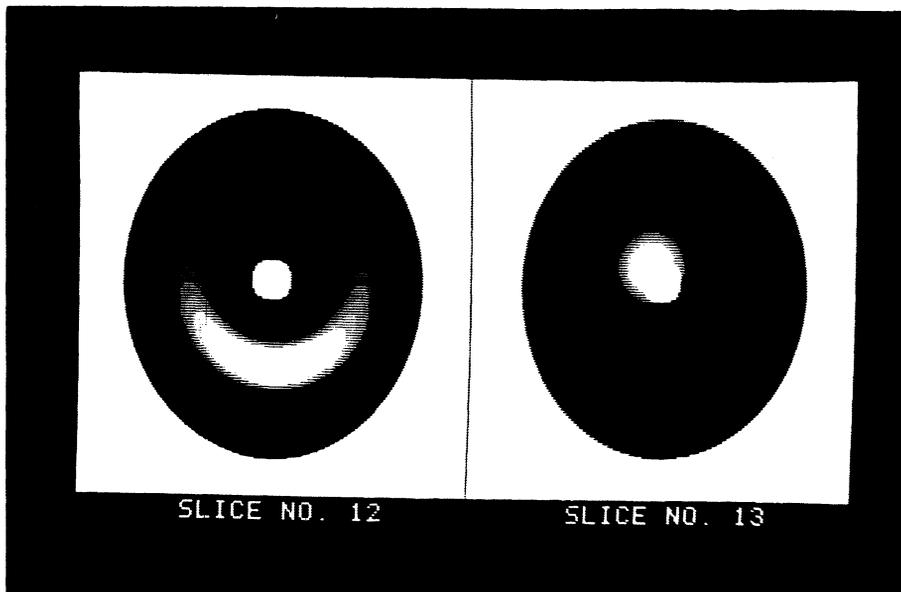


Figure 10c. simulated CT slices.

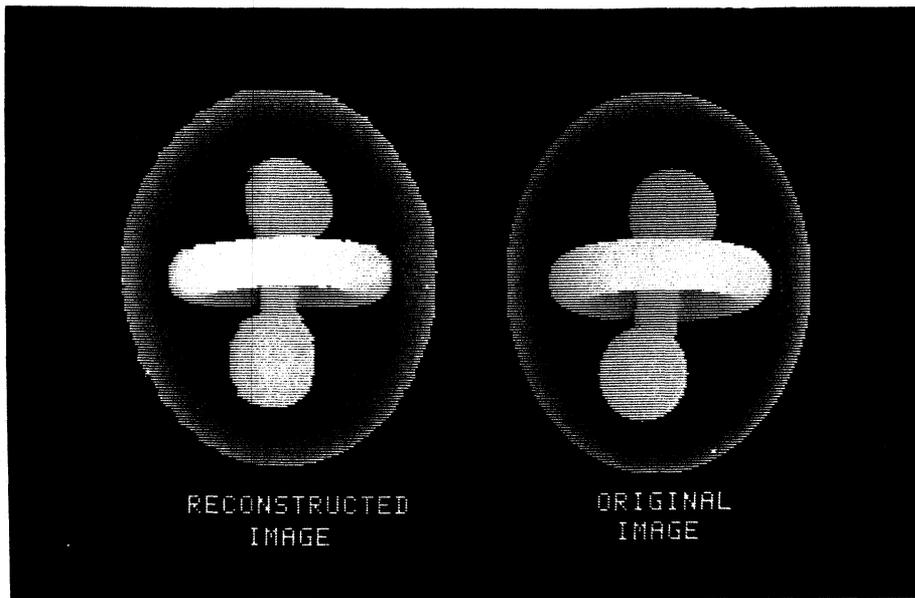


Figure 11a. The reconstructed and the original objects.

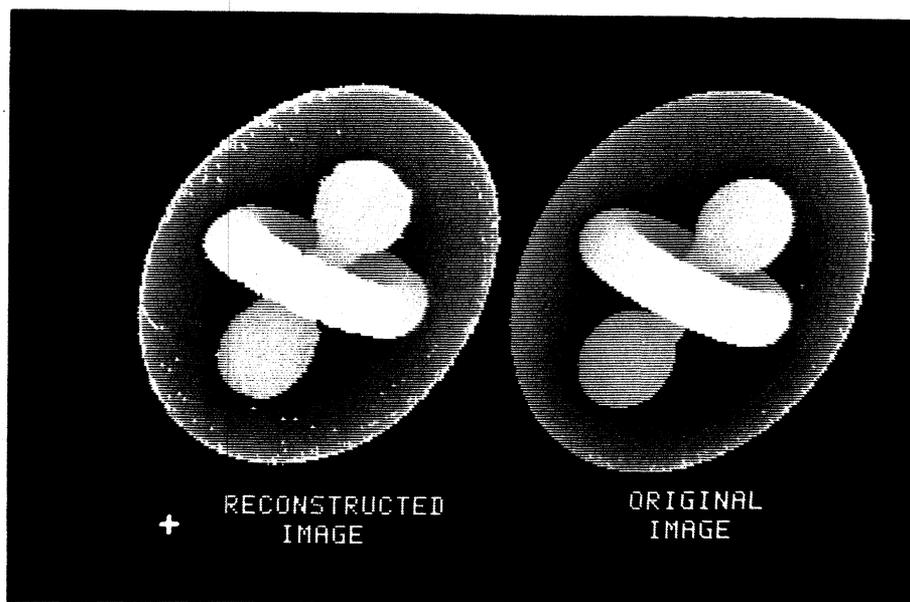


Figure 11b. The reconstructed and the original objects.

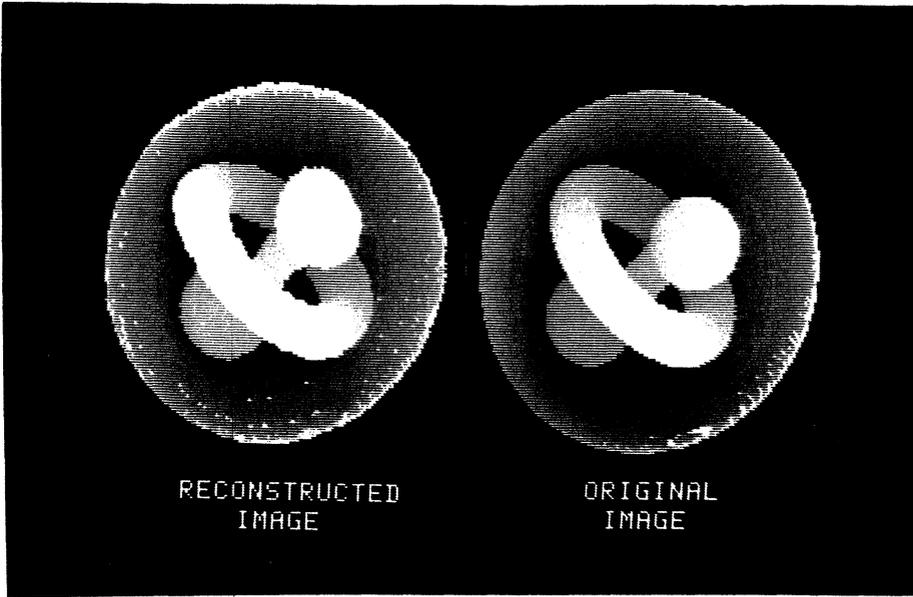


Figure 11c. The reconstructed and the original objects.

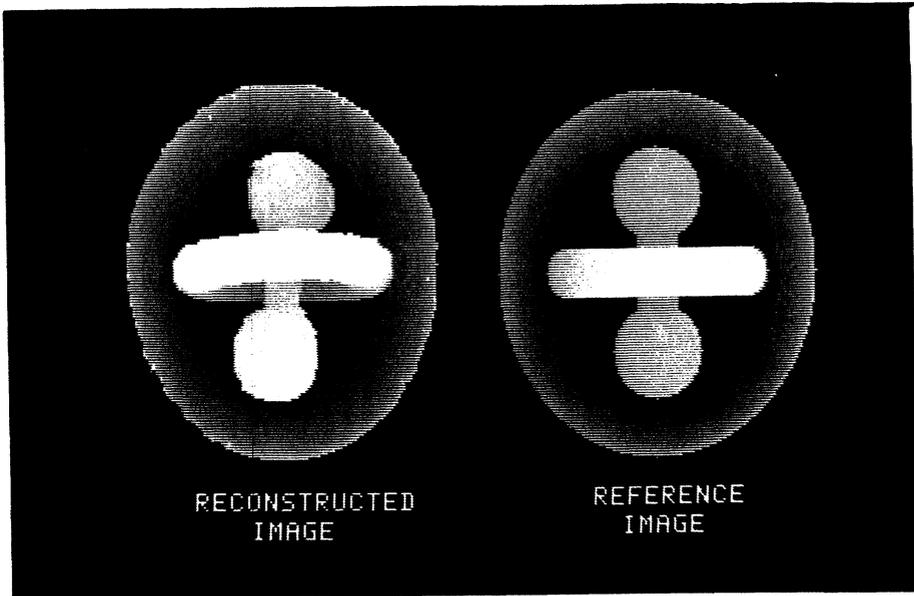


Figure 12a. The reconstructed and the reference objects.

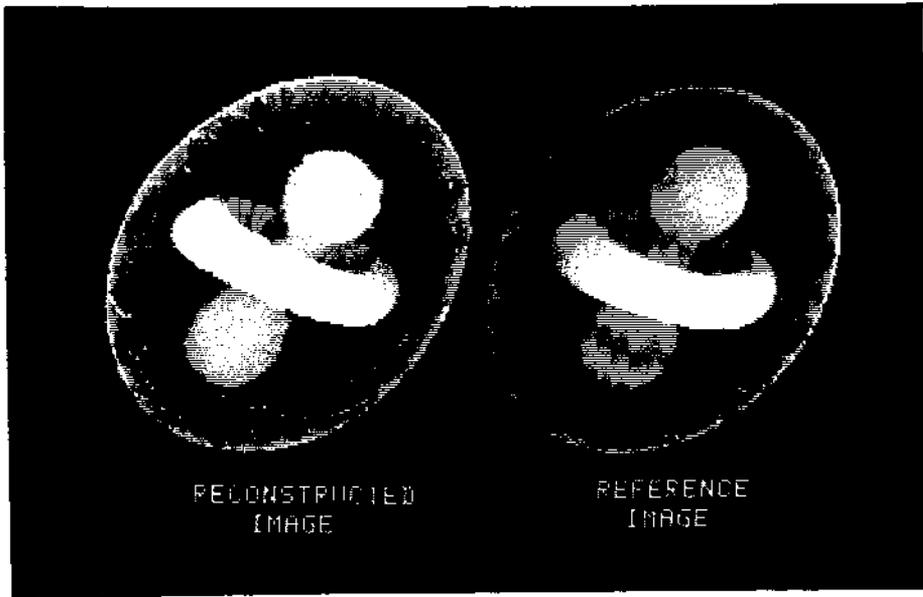


Figure 12b. The reconstructed and the reference objects.

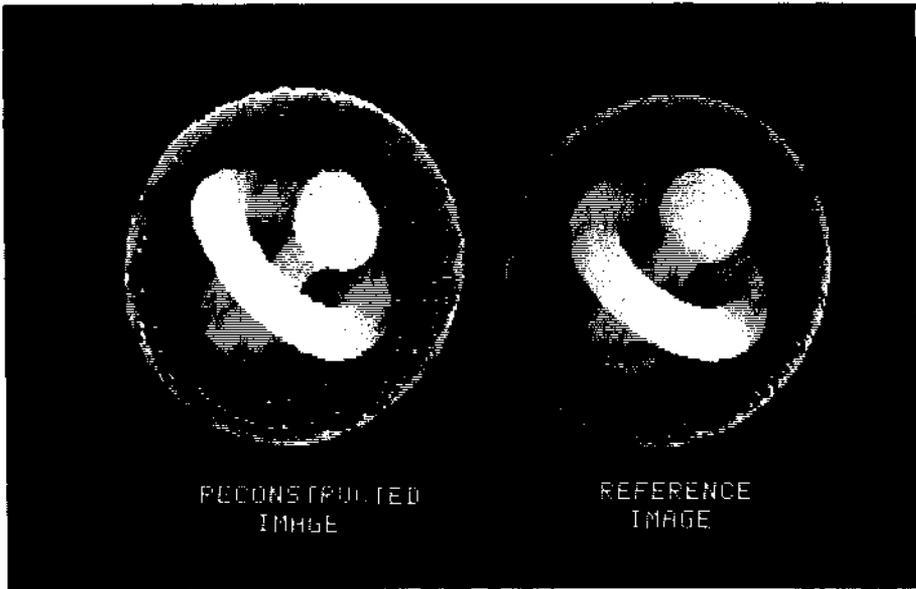


Figure 12c. The reconstructed and the reference objects.

It is clear that the shape of the reconstructed object falls between the shape of the original object and the reference object* The shapes of the reconstructed ellipsoid and the reconstructed torous are quite close to the shapes of their corresponding objects in the original image but larger errors occurred in the reconstructed barbell* These errors are due to the larger deformations required to translate the barbell through the torous*

The final example is the registration of an image obtained by stacking fifteen CT slices with the anatomy atlas of the brain* The atlas used in this example was constructed from thirty slices with a thickness of 4 mm each* Four of the slices in this atlas are shown in figures 5a, 6a, 7a and 8a* Six views of the atlas as a three dimensional semi-transparent object are shown in figures 13a-13c* Four views of the reconstructed brain are shown in figures 14a and 14b*

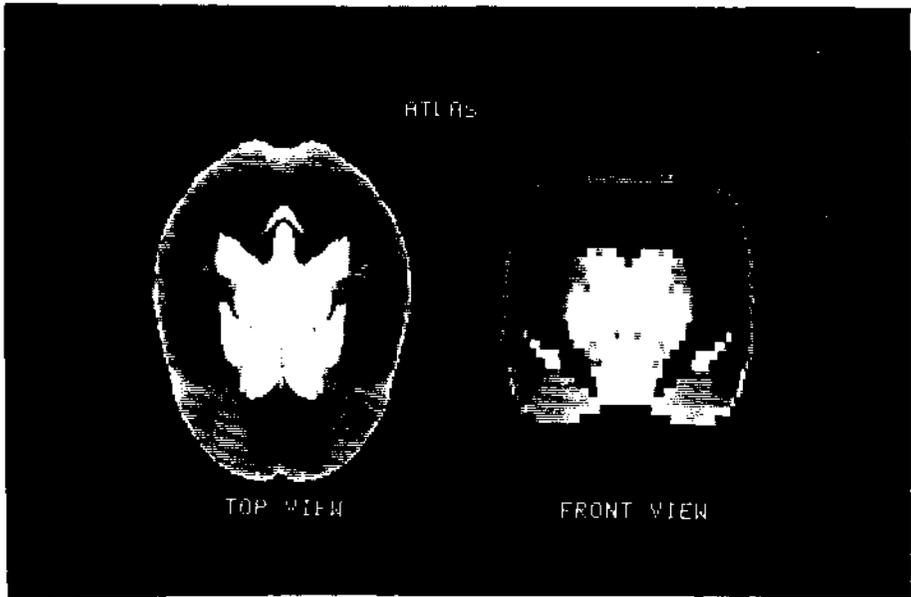


Figure 13a. Three dimensional atlas*.

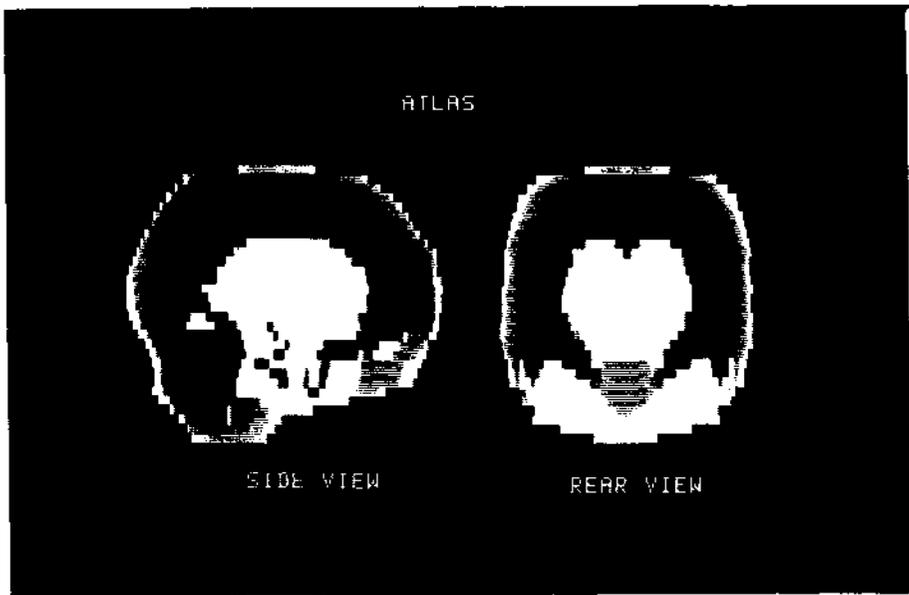


Figure 13b. Three dimensional atlas.

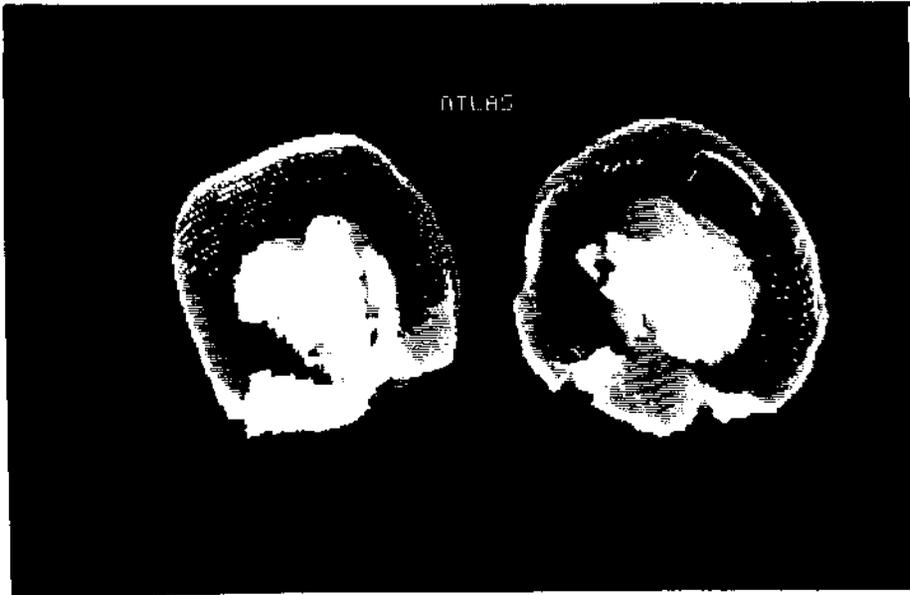


Figure 13c. Three dimensional atlas.

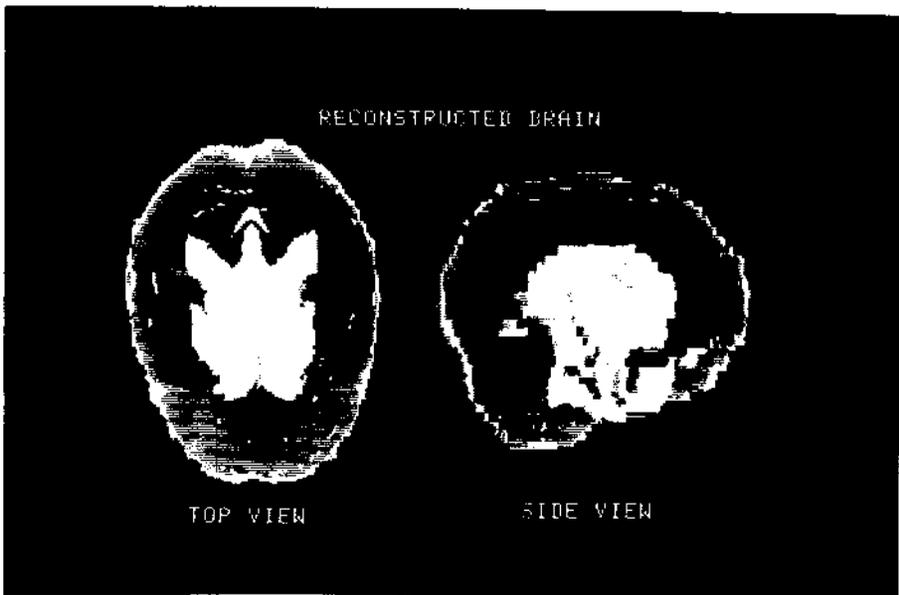


Figure 14a. The reconstructed brain

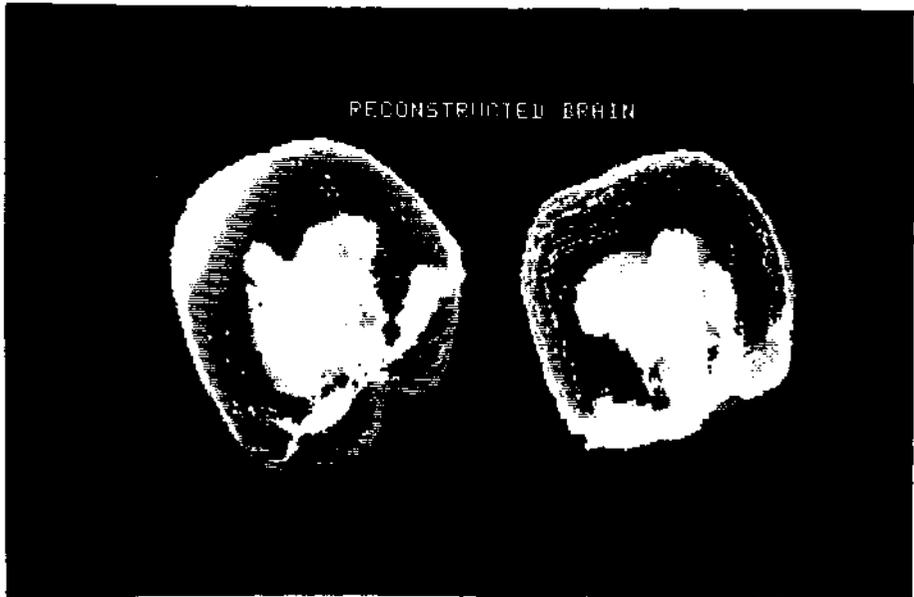


Figure 14b. The reconstructed brain.

Figure 15 shows the top view of the reconstructed brain (left hand side of the figure) and the top view of the atlas (right hand side). By comparing these two views it is possible to see that the reconstructed brain is rotated (anti-clockwise) with respect to the atlas and that the left side of the reconstructed brain is closer to the viewer than the right side (the left side is brighter than the right side). Unfortunately, it was not possible to judge the accuracy of the reconstructed brain since the CT images were taken from a living person.

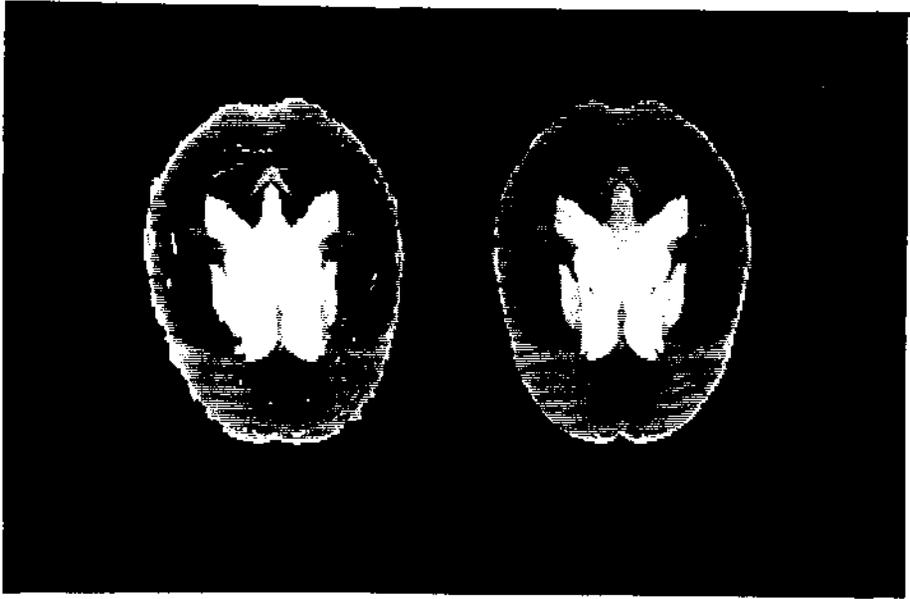


Figure 15. The reconstructed brain and the atlas.

CHAPTER 9 CONCLUSION

9*1 Summary

The work described in this dissertation was motivated by the need to locate objects and their boundaries in three dimensional images obtained by stacking successive CT images* This task is difficult because of the complexity of the anatomic structure and also because of the poor quality of the CT images. Automatic methods that do not use external sources of knowledge had only limited success in the past. The use of external knowledge usually requires sophisticated programs and complex knowledge representation techniques* While in the long run this approach is very promising, in the short run there are very few results*

In this work, the external knowledge was represented simply by another three dimensional image - the atlas image* By registering the CT image with the atlas image, a mapping between the two is obtained* This mapping enables us to superimpose the boundaries from the atlas on the CT image and use them as approximations for the true boundaries*

The major problem that we had to solve was that of finding the mapping. When the differences between the two images cannot be accounted for by a rigid transformation only, there is some freedom in selecting the mapping. The mapping is evaluated by the similarity of the transformed image to the other image and by the amount of deformation. By defining a cost function that contains these two values we can define an optimal mapping.

To obtain this optimal mapping we need tools to measure the similarity under different deformations and the deformations themselves. Given these, an efficient procedure to locate the minimum cost is necessary. When the image is treated as an array of points, even a simple version of this problem is extremely difficult to solve [FISCHLER 1973]. If instead we treat it as a continuum we can use the theory of elasticity for the problem.

Taking the similarity as the potential function from which forces are derived to deform the image, we can use Navier's equations to solve the problem even without giving an explicit expression for the deformation. These equations are solved by an iterative method on a grid of points. The mapping is described by the placements of the grid points. Approximate values of these can be used in computing the similarity under deformation.

The similarity is measured by the cross correlation of the image functions in two regions. This operation, which is repeated many times during any registration process, consumes the largest part of the computation time. To save time we have expressed the image function in a region by its projections on a set of orthogonal functions. Less than a dozen projections are usually sufficient for the process, he computation time. To save time we have expressed the image function in a region by its projections on a set of orthogonal functions. Less than a dozen projections are usually sufficient for the process, instead of the more than several hundred pixels that are contained in the region. Using the solid spherical harmonics as the set of orthogonal functions, it became possible to compute these projections in a simple and fast way under different rotations from the original set of projections.

Our registration method is also very useful in finding the mapping between two images when a set of pairs of corresponding points are available. Not only is the mapping obtained by our method optimal, but it is also possible to vary the ratio between the deformation and the fitting of the data. This mapping can also take into account different error estimates for each observation and these errors can have asymmetric probability distribution.

9.2 Relationship To Other Methods

Most of other registration methods contain two parts, one dealing with with the local aspects of the matching and the other with the global ones* The task of the first part is to locate a set of corresponding points that will be used by the second part to compute the mapping. The task of the second (or global) part is complicated by the errors that the first part makes. These errors are the results of distortion in gray level, geometrical deformations, noise in the images and insufficient details for unique identification.

Most methods rely on the assumption that the mapping is continuous in handling these errors* The assumption is used either by fitting a low degree polynomial to the data or by averaging the displacements of adjacent points. In both cases, no attempt is made to estimate the errors. Only Fischler's method [FISCHLER 1973] uses a cost function containing both the matching goodness as an estimation of the error and the deformation.

Cooperation between the two parts, that is the guidance of the local matching process, is minimal, while it is clear that the similarity between the images is a function of the computed mapping, only Burr [BURR 1979] uses an iterative method to improve the correlation measure. He

does it by deforming the two images after every iteration step.

Our registration method is a generalization of Fischler's, Burr's and Tobler's [TOBLER 1978] methods. We have used a cost function similar to that of Fischler, a grid to represent the mapping as Tobler does, and the similarity is computed in every iteration based on the mapping obtained in the previous iteration. The contribution of this work is not just in combining these ideas into an efficient method, but also in significantly improving each one of them.

The model of a spring chain used by Fischler and Elschlager is too simple for an effective measure of deformation. A triangulated network of springs would be much better but their method cannot handle it in a polynomial time. The deformation part of our cost function is based on a model of a continuous solid. Even the finite grid used for the iterative solution contains (in the three dimensional case) several thousands of points compared with less than a dozen in their model. The mapping obtained by our method is optimal and theirs only in the case of the linear chain.

Burr has to deform the two images in each iteration step. Our method has only to deform the projections for the points of one grid whose total number is much less than the number of pixels in the image. The use of these projections also enable us to compute the similarity between two regions by only a few operations and therefore we have a significant speed advantage over methods using direct method for computing the cross correlations.

Tobler's method does not use the error estimates that are associated with each pair of observation and it is also not possible to vary the stiffness of the grid. In our method the stiffness of the grid is easily changed by varying one or two elastic constants. In that respect it is similar to varying the degree of the multivariate polynomial that represents the mapping. The error in each observation is estimated from the cross correlation function and the resulting mapping is optimal in the same sense as the least squares method.

9.3 Suggestions For Future Research

Three directions for further research are suggested here. The first deals with improving the atlas so that the use of our method for locating object boundaries in three dimensional CT images can be tested more thoroughly. The

second is the improvement of the method so that larger deformations could be handled. The third direction is the study of other possible applications of the registration method.

The main difficulty that we had in the application of the registration method is due to the lack of a good atlas. The construction of a good atlas turn out to be more difficult than we had expected. In our opinion a good atlas should contain two versions. One which resembles as much as possible a CT image in its density levels. The other version should contain the structure of the object boundaries. The representation of this version should enable reslicing of the atlas through curved surfaces. These slices will be superimposed on the CT image. Two versions are required because certain parts of the anatomy ,for example the skin that encloses the brain, have high enough density to appear in the image while their structure cannot be represented easily and is not required.

The assumption underlying our registration method is that the deformation part of the mapping is small. This assumption is not always justified, but the main reason for making it was its simplicity. If this assumption is not used, the equilibrium equations become nonlinear and therefore much more difficult to solve. Nevertheless, nonlinear (or finite) deformations are part of the theory of

elasticity and such problems are treated in the literature [GREEN 1968] and [GREEN 1970]).

An important application which could benefit from this method is the processing of stereo images and in particular the generation of topographic maps from aerial images* The large size of of these images and their high resolution make their processing by other methods difficult* Using only a small number of projections for correlation our method can be fast enough for this purpose*

Another application that could perhaps be developed following a method similar to the elastic matching, is the construction of smooth surfaces from a given set of points* By modeling the surface as an elastic sheet that is pulled by the points, an optimal fitting could be achieved*

BIBLIOGRAPHY

[ARTZY 1981]

Artzy, E., Frieder, G. and Herman, G.T., "The Theor Design, Implementation and Evaluation of a Three Dimension Surface Detection Algorithm", Comp. Graph. & Image proces 15, No. 1, Jan. 1981, pp. 1-24.

[BARNEA 1972]

Barnea, D.I., and Silverman, H.E., "A Class Algorithms for Fast Digital Image Registration", IEEE Tran Comput. C-21, no. 2, Feb. 1972, pp. 179-186.

[BURR 1979]

Burr, D.J., "A Dynamic Model for Image Registration IEEE Computer Society Conference on Pattern Recognition and Image Processing, Aug. 1979, Chicago Ill., pp. 17-24.

[CHRISTIANSEN 1978]

Christiansen, N.H. and Sedderberg, T.W., "Conversion Complex Line Definitions into Polygonal Element Mosaic Computer Graphics 12, No. 3, Aug 1978, pp. 187-192

[COOK 1980]

Cook, P., "Three Dimensional Reconstruction From Serial Sections For Medical Applications", Ph.D. Dissertaio University of Missouri - Columbia, 1980.

[DAVIS 1975]

Davis, L.S., "A Survey of Edge Detection Techniques Comp. Graph. & Image Proc. 4, No. 3, Sept. 1975, p 248-270.

[FISCHLER 1973]

Fischler, M.A., and Elschlager, R.A. "T Representation and Matching of Pictorial Structure", IE Trans. Comput. C-22, no. 1, Jan. 1973, pp. 67-92.

[FUCHS 1977]

Fuchs, H., Kedem, Z.M. and Uselton, S.P. "Optimal Surface reconstruction from Planar Contours", Comm. ACM 20, 1977, pp. 693-702.

[GLENN 1977]

Glenn, W.V., Davis, K.R., Larsen, G.N. and Dwyer, S.J., "Alternative Display Formats for Computed Tomography Data", Current Concepts in Radiology, Vol. 3, Potchen, E.J. ed. C. V. Mosby Co., St. Louis, Missouri 1977.

[GREEN 1968]

Green A.E. and Zerna W. "Theoretical Elasticity", Clarendon Press, Oxford, 1968

[GREEN 1970]

Green A.E. and Adkins J.E., "Large Elastic Deformations", Clarendon Press, Oxford, 1968

[GREENLEAF 1970]

Greenleaf, J.F., Tu, J.S. and Wood, E.H. "Computer Generated Three Dimensional Oscilloscopic Images and Associated Techniques for Display and Study of the Spatial Distribution of Pulmonary Blood Flow", IEEE Trans. Nucl. Sci. NS-17, June 1970, pp. 353-359.

[GRIMSON 1980]

Grimson, W.E.L., "A Computer Implementation of a Theory of Human Stereo Vision", A.I. Memo No. 565, M.I.T. Artificial Intelligence Lab. 1980.

[HANNAH 1974]

Hannah, M. J., "Computer Matching of Areas in Stereo Images" AIM-239, Computer Science Dept., Stanford Univ., July 1974.

[HERMAN 1979]

Herman, G.T., Ed. "Image Reconstruction from Projections: Implementation and Applications", Springer-Verlag, Berlin New-York, 1979.

[HOBSON 1931]

Hobson, E.W., "The Theory of Spherical and Ellipsoidal Harmonics", Cambridge University Press, 1931.

[HUECKEL 1971]

Hueckel, M.H., "An Operator which Locates Edges in Digital Pictures" J. Assoc. Comput. Mach. 18, no. 1, Jan. 1971, pp. 113-125.

[HUECKEL 1973]

Hueckel, M.H., "A Local Visual Operator which Recognizes Edges and Lines", J. Assoc. Comput. Mach., 20, no. 4, Oct. 1973, pp. 634-637.

[KEPPEL 1975]

Keppel, E., "Approximating Complex Surfaces By Triangulation Of Contour Lines", IBM J. Res. Develop. 19, Jan. 1975, pp. 2-11.

[LIU 1977]

Liu, H.K., "Two and Three Dimensional Boundary Detection", Comp. Graph. & Image Process. 6, 1977, pp. 123-134

[MARR 1979]

Marr, D., and Poggio, T. "A Computational theory of Human Stereo Vision", Proc. R. Soc. Lond. B-204, 1979, pp. 301-328.

[O'ROUKE 1980]

O'Rourke, J., "Image Analysis of Human Motion", Ph.D. Dissertation, Computer and Information Science, U. of PA., 1980.

[PRATT 1974]

Pratt, W.K., "Correlation Techniques of Image Registration", IEEE Trans. Aerosp. Electron. Syst. AES-10, no. 3, May 1974, pp. 353-358.

[RHODES 1979]

Rhodes, M.L., ^{ff}"An Algorithmic Approach to Controlling Search in Three Dimensional Image Data", SIGGRAPH '79 Proceedings, Chicago, 111., Aug• 1979, pp. 134-142.

[ROSENFELD 1976]

Rosenfeld, A., and Kak, A.C., "Digital Picture Processing", Academic Press, New York, 1976.

[SADJADI 1978]

Sadjadi, F., and Hall, E.L., "Invariant Moments for Scene Analysis", Proc. IEEE Conf. Pattern Recognit Image Process., Chicago 111., May 1978, pp. 181-187.

[SOKOLNIKOFF 1956]

Sokolnikoff, I.S., "Mathematical Theory of Elasticity" McGraw-Hill Book Co., New York, 1956.

[SUNGUROFF 1978]

Sunguroff, A.S. and Greenberg, D., "Computer generated Images for Medical Application", SIGGRAPH '78 Proceedings, Atlanta Ga., Aug. 1978, pp. 196-202.

[TOBLER 1977]

Tobler, W.R., "Bidimensional Regression", Geography Dept., Univ. of Calif., Santa Barbara, 1977.

[TOBLER 1978]

Tobler, W.R., "Comparing Figures by Regression" Computer Graphics 12, No. 3, Aug 1978, pp. 193-195.

[WIDROW 1973]

Widrow, B., "The 'Rubber Mask' Technique", Pattern Recognition 5, no. 3, Sept. 1973, pp. 175-211.

[WONG 1977]

Wong, R.Y., "Image Sensor Transformations", IEEE Trans. Syst. Man Cybr. SMC-7, no. 12, Dec. 1977, pp. 836-841.

ONG 1978]

Wong, R.Y., and Hall, E.L., "Sequential Hierarchical
ene Matching", IEEE Trans. Comput. C-27, no. 4, Apr. 1978,
. 359-365.

OUNG and GREGORY 1973]

Young, D.M., and Gregory, R.T., "A Survey of Numerical
thematics", Addison-Wesley, Reading Mass. 1973.

UCKER 1976]

Zucker, S.W., "Region Growing: Childhood and
olescence", Computer Graphics & Image Processing 5, 1976,
. 382-399.

INDEX

- Artzy E. 16
- Atlas 22,23,106,111,131,134,136,141-142
- Autocorrelatio 105
- Base functions 65-70
- Beltrami-Michel equations 52
- Barnea D.I. 27
- Boundary conditions 10,51,96-98
- Brain 13,15,16,22,131,134
- Burr D.J. 34,139-141
- Cauchy A. 9
- Center of mass 74
- Christiansen N«H« 19
- Compatability 51,52
- Continuity assumption 40
- Convolution theorem 64
- Cook >. 19
- Cost function 1,2,6,45,137
- Cross correlation 9,26,54-65,83,102,104,105,110,138,141
 - normalized 60
- Davis L.S. 17
- Deformation 2-6,10,39-51,54,56,85,100,106,138,143
- Dilationless stretch 48
- Dirichlet problem 52
- Elastic
 - constatnts 6,10,82,108-110,141
 - registraion 2,3,4,143
- Elasticity 9,10,83,137
- Equilibrium
 - state 5,6,49,76
 - equations 48,76,142
- External
 - forces 5,7,48,51,52,80,82-90,93,104,106,108
 - knowledge 20,136
- Finite differences 10,76-79,100,108
- Fischler M.A. 33,34,137-140
- Fourier transform 64
 - fast 65
- Fuchs H. 18
- Gauss's theorem 49
- Gauss-Siedel method 80,81,95
- Glenn W.V. 14
- Global mapping 96,102-104

Green G. 9
 Green A.E. 143
 Gregory R.T. 81
 Grimson W.E.L. 32
 Hannah M.J. 28,30
 Herman G.T. 12
 Hilbert space 59
 Hirarchical search 31
 Hobson E.W. 67
 Homogeneous deformation 47
 Hooke's law 50,109
 Hueckel M.H. 10,69
 Inertia tensor 73
 Intrinsic
 coordinates 73
 orientation 73
 Inverse mapping 112
 Invariant
 features 30,66,72-75
 moments 29
 Jaccobi's method 80,95,108
 Keppel E. 17
 Lamé's constants 50
 Lagrange's model of elasticity 84
 Laplacian 28,78
 Laplace's equation 97
 Least squares approximation 87,103-106,141
 Legendre's polynomials 68
 Liu H.K 17
 Markov process
 Marr D. 29,32
 Mesh size 10,77,100-101,104,112
 Neumann problem 52,96
 Navier 9
 Navier's equations 48,76,137,142
 Orthogonal
 base 99
 functions 7,9,59-72,138
 O'Rourke J. 21
 Optimal
 mapping 1,2,5,91
 registration 1,111
 Plastic registraion 2
 Poggio T. 29
 Polynomial wrapping 35
 Poisson's ratio 51,53
 Pratt W.K. 27
 Principal
 axes 46
 extensions 46
 Projections 7,12,59-72,99-102,104,138,141

lection 40
ression
 bidimensional 37,91
 three-dimensional 90-95
des M.L. 16
enfeld A. 25,33
ation 1,4,41,65,66,70-75,88,103,107,124
ber mask 33
ljadi F. 29,66
nt-Venant principle 53
ilarity 2,4,7,26,29,48,54-57,76,83-85,100,106,138,141
ple expansion 47
 shear 47
olnikoff I.S. 52,53
erical harmonics 67-69,138
ring network 33,34,140
nguroff A.S. 18
atistical correlation 28
mplate matching 25,26
bler W.R. 36,37,91,140-141
ansformation
 affine 4,41,42,44,65,71
 global 4,6,45
 local 44,45
 projective 4
 rigid 39,41,124
anslation 1,4,9,41,65-66,103,124
ak dominant diagonal condition 82
ldrow B. 33
ng T.Y. 31,35
ung D.M. 81
ung's modulus 51
ero crossings 101